

Computational Dynamics

Spring 2023

Lecture Notes

Jaroslav Borovička[†]

March 13, 2023

[†]E-mail: jaroslav.borovicka@nyu.edu. All errors are mine.

Contents

1	Markov chains and asset valuation	3
1.1	Markov chains	4
1.1.1	Perron–Frobenius theorem	5
1.1.2	Conditional, unconditional, and stationary distributions	5
1.1.3	Forecasting and conditional expectations	9
1.1.4	Invariant functions and ergodicity	10
1.1.5	Simulation and likelihood estimation	14
1.1.6	Method of moments estimation	15
1.1.7	Hidden Markov models	17
1.1.8	Continuous-state Markov chain	20
1.1.9	Example: Asset valuation	21
1.1.10	Example: Wages in a lake model	24
1.2	Investors’ consumption-saving decision problem	26
1.2.1	Subjective beliefs	29
1.2.2	Transversality conditions vs No-Ponzi conditions	30
1.3	Equilibrium in endowment economies	32
1.3.1	Equilibrium asset prices	34
1.3.2	Risk premia as compensation for risk	36
1.4	Testing Euler equations using generalized method moments	37
1.4.1	Role of rational expectations in testing Euler equations	40
1.4.2	Generalized method of moments	40
1.5	Valuation in finite-state Markov chain economies	43
1.5.1	Asset pricing formulas	45
1.5.2	Equity premium puzzle in the Mehra–Prescott economy	48
1.6	Absence of arbitrage and equilibrium asset prices	51
2	Value function iteration in search problems	55
2.1	A simple model of intertemporal search	55

2.1.1	Information sets, beliefs, and strategies	56
2.1.2	Recursive formulation	57
2.1.3	Principle of optimality	59
2.1.4	Characterizing optimal policy	60
2.1.5	Mean-preserving spreads	62
2.1.6	Effects of mean-preserving spread in the McCall (1970) model	65
2.1.7	Welfare	65
2.2	Extensions	67
2.2.1	Option to quit	67
2.2.2	Firing	68
2.2.3	Finite horizon	70
2.2.4	Backward induction	71
2.2.5	Welfare in a model with exogenous transitions	72
2.2.6	Further extensions	74
2.3	Equilibrium wage distribution	74
2.4	Numerical implementation	75
2.4.1	Iteration on the reservation wage	75
2.4.2	Value function iteration	76
2.4.3	Evaluating the expectations operator	78
2.4.4	Gaussian quadrature	79
3	Perturbation methods and linear state space models	93
3.1	Linear vector autoregressions	94
3.1.1	First and second moments	96
3.1.2	Impulse response functions	99
3.1.3	Applications	101
3.1.4	Likelihood of the linear vector autoregressive process	105
3.2	Perturbation methods	106
3.2.1	Taylor's theorem	106
3.2.2	Series expansion method	107
3.2.3	Campbell–Shiller decomposition of returns	109
3.2.4	Linear approximation of valuation ratios	112
3.2.5	Linear-quadratic approximation of stochastic control models	115
3.2.6	Solving linear equilibrium models	117
3.2.7	VAR estimation and identification of structural shocks	118
3.2.8	A cautious note on the accuracy of perturbation approximations	120

4	Filtering and learning	121
4.1	Kalman filter	121
4.2	Applications	124
4.2.1	Muth's example	124
4.2.2	Jovanovic's model	126
4.2.3	Innovations representation and VARs	127
4.2.4	Relationship to linear-quadratic dynamic programming	128
4.3	Estimation	130
4.3.1	Estimation of the hidden state space model	130
4.3.2	Frequentist and Bayesian inference	132
5	Finite difference methods in derivative pricing	133
5.1	Brownian motion	134
5.1.1	Motivating examples	134
5.1.2	Definition and properties	138
5.1.3	Variation of a stochastic process	139
5.2	Itô calculus	141
5.2.1	Construction of Itô integral	142
5.2.2	Itô processes	146
5.2.3	Itô's lemma	147
5.2.4	Examples	149
5.2.5	Feynman–Kac formula	151
5.3	The Black–Scholes model	152
5.3.1	Market structure	153
5.3.2	Derivative securities	154
5.3.3	Replication argument	155
5.3.4	Analytical solution	156
5.3.5	Comparison with Feynman–Kac formula and risk-neutral valuation	157
5.3.6	Extensions and further remarks	158
5.3.7	Numerical implementation	159
5.4	Difference schemes for parabolic PDEs	160
5.4.1	Mapping of the problem to the Black–Scholes model	161
5.4.2	Grid construction and approximation with differences	162
5.4.3	Explicit solution scheme	164
5.4.4	Implicit solution scheme	166
5.4.5	Upwind scheme for the choice of first difference	167
5.4.6	Approximation errors	169

5.4.7	Nonlinear equations and steady states	170
5.5	Binomial trees	172
6	Q-learning in decision problems	177
6.1	Monte Carlo methods	178
6.2	Q-learning	179
6.2.1	Adding decision-making	180
6.2.2	Interpretation	181
6.2.3	Implementation details	182
6.3	Q-learning in the worker search problem	183
6.3.1	Implementing Q-learning	184
6.3.2	Implementation details	186
6.3.3	Results	187
6.4	Summary	190
A	Elements of probability theory	193
A.1	States, paths, partial histories and probabilities	193
A.2	Filtration	195
A.3	Random variables and stochastic processes	196
A.4	The Radon–Nikodým derivative	198
A.5	Expectations	199
A.6	Martingales	201
A.7	The Markov property	201
A.8	Macroeconomic models and parameters	203

Introduction

These lecture notes cover the Computational Dynamics course taught in the Master of Science in Quantitative Economics (MSQE) program at the Department of Economics at NYU. The material covers a series of computational methods applied to economic problems in macroeconomics and asset pricing.

Substantial parts of the theoretical material are covered in [Ljungqvist and Sargent \(2018\)](#). A more in-depth coverage of dynamic optimization is provided in [Stokey et al. \(1989\)](#). A range of related topics in asset pricing theory is discussed in [Duffie \(2001\)](#) and [Cochrane \(2005\)](#). Elements of applied modeling of dynamic stochastic general equilibrium economies can be found in [Galí \(2008\)](#).

Detailed analysis of numerical methods with applications to economics is provided in [Judd \(1998\)](#). Numerical methods for the solution of differential equations can be found in [Holmes \(2007\)](#). A variety of approaches in machine learning are contained in [Bishop \(2006\)](#), [Goodfellow et al. \(2016\)](#), or [Sutton and Barto \(2018\)](#).

An excellent source of supporting material is the QuantEcon website developed by John Stachurski, Tom Sargent and other helpful contributors:

<https://quantecon.org/>.

Specifically, we make use of the lectures that focus on Python implementation of numerical algorithms and economic models.

- *Python Programming for Economics and Finance* covers instructions on how to install the Python environment and make it operational with QuantEcon, together with elementary Python programming:

<https://python-programming.quantecon.org>.

- *Quantitative Economics with Python* is a course that covers basic economic modeling:

<https://python.quantecon.org>.

- *Advanced Quantitative Economics with Python* discusses more advanced economic topics:

<https://python-advanced.quantecon.org>.

Python and Matlab implementation of the computational methods discussed in these lecture notes is also available on the associated GitHub website

<https://github.com/jborovicka/nyu-computational-dynamics>.

Chapter 1

Markov chains and asset valuation

Textbook: Ljungqvist and Sargent (2018), Chapters 2 (Sections 2.2–2.3, Markov chains), Chapters 13 and 14 (asset pricing theory and empirics). Cochrane (2005), Chapters 10–13 (application of GMM in asset pricing)

Generalized method of moments: Hansen (1982), Hansen (2008).

Asset pricing applications: Lucas (1978), Hansen and Singleton (1982, 1983), Mehra and Prescott (1985).

QuantEcon: Quantitative Economics with Python: Topic 3 (linear algebra), Topic 25 (finite Markov chains), Topics 74–76 (asset pricing applications in finite state models). **Advanced Quantitative Economics with Python:** Topics 34–35 (more advanced asset pricing applications).

Asset pricing seeks to explain time-series and cross-sectional patterns in the valuation of financial assets. Asset valuations are determined by three key factors—risk embedded in cash flows the valued asset generate, investors’ attitudes to this risk, and investors’ ability to insure these risks using available markets or self-insurance. These factors interact and jointly determine the equilibrium in the macroeconomy.

In order to embed the study of asset prices in a tractable environment, we introduce stochastic processes called Markov chains. A Markov chain is stochastic process that takes values on a finite state space and satisfies the Markov property, which states that the conditional distribution of future realizations of the process conditional on information available up to a given time period is identical to the conditional distribution that conditions only on the current value of the stochastic process in that period. This is a critical assumption that will make the characterization of investors’ optimal decisions and computation of expected values of cash flows tractable.

We start by introducing the Markov chain model in Section 1.1. In Section 1.2, we introduce a consumption-saving problem of a utility-maximizing investor. From this model, we derive conditions for investors’ optimal investment behavior in the form of Euler equations. These Euler equations constitute testable restrictions for the joint behavior of cash flows, asset prices, and investors’ marginal rate of substitution.

Since the testable restrictions from investors’ optimizing behavior come in the form of

moment conditions, we can test them using the generalized method of moments (GMM). We discuss the applications in Section 1.4.

The GMM tests do not require a full specification of the model environment nor the stochastic structure imposed by a Markov chain. As a next step, we embed investors' consumption-saving problem in a representative agent endowment economy in the spirit of Lucas (1978). We impose Markov chain structure and fully solve for the equilibrium asset prices in a replication of the model of Mehra and Prescott (1985).

Finally, in Section 1.6, we briefly connect the existence and structure of equilibrium asset prices to the concept of absence of arbitrage opportunities.

Foundations of results from probability theory that encompass the Markov chain model and that will also be useful in subsequent chapters can be found in Appendix A.

1.1 Markov chains

A special type of a Markov process is a **Markov chain**. The structure of transition probabilities of a Markov chain allows to model arbitrary nonlinear dynamics, at the expense of having to restrict the model to a relatively small number of discrete states to keep it computationally tractable. We focus here on time-invariant, finite-state, discrete-time Markov chains, constructed in the following way.

Definition 1.1. An n -state *time-invariant Markov chain* is defined by the following triplet of objects:

- an n -dimensional state space consisting of vectors e_i , $i = 1, \dots, n$, where e_i is the coordinate vector with a single 1 in i -th entry;
- an $n \times n$ transition matrix \mathbf{P} (defining the transition probability) such that

$$\mathbf{P}_{ij} = P(x_{t+1} = e_j \mid x_t = e_i);$$

- an $n \times 1$ vector π_0 that determines the initial distribution of the state

$$\pi_{0i} = P(x_0 = e_i).$$

Each row i of the transition matrix \mathbf{P} determines the distribution of the state next period conditional on state e_i today. In order for π_0 and \mathbf{P} to define a valid probability measure, we require

$$\sum_{j=1}^n \pi_{0j} = 1, \quad \sum_{j=1}^n \mathbf{P}_{ij} = 1 \quad \forall i = 1, \dots, n.$$

Time-invariance in the context of the Markov chain means that \mathbf{P} does not depend on t .

1.1.1 Perron–Frobenius theorem

The following central result from linear algebra will play an important role in the characterization of asymptotic dynamics on Markov chains.

Proposition 1.1 (Perron–Frobenius). *Let \mathbf{A} be an $n \times n$ matrix with strictly positive entries, $\mathbf{A}_{ij} > 0, \forall i, j$. Then the following is true.*

- *There is a positive real number r , called the **Perron root** (or **Perron–Frobenius eigenvalue**) such that r is an eigenvalue of \mathbf{A} , and any other eigenvalue λ (possibly complex) is strictly smaller than r in absolute value, $|\lambda| < r$. Thus the spectral radius is $\rho(\mathbf{A}) = r$.*
- *The Perron–Frobenius eigenvalue is simple: r is a simple root of the **characteristic polynomial** of \mathbf{A} , the polynomial given by $\det(\mathbf{A} - \lambda\mathbf{I})$ as a function of λ . Consequently, the **eigenspace** associated to r is one-dimensional, there exists (up to scale) only one nonzero **eigenvector** v such that $\mathbf{A}v = rv$.*
- *The eigenvector v can be normalized to have all elements strictly positive, $v_j > 0, \forall j = 1, \dots, n$. There are no other non-negative real eigenvectors associated with different eigenvalues.*

The assumption of strict positivity of all elements of the matrix, imposed on the transition matrix \mathbf{P} , implies that it is possible to transit from any state to any state in one period with a strictly positive probability, a property known as a **mixing condition**, or **irreducibility**.

1.1.2 Conditional, unconditional, and stationary distributions

The transition matrix \mathbf{P} represents 1-period ahead conditional probabilities of individual states. Conditional n -period ahead probabilities are given by the elements of the matrix \mathbf{P}^n :

$$(\mathbf{P}^n)_{ij} = P(x_{t+n} = e_j \mid x_t = e_i).$$

The **unconditional distribution** of the initial state is $\pi_0 = P(x_0)$. Using forward iteration on the distribution using the transition matrix yields the unconditional distribution at time $t = 1$

$$\pi'_1 = P(x_1) = \pi'_0 \mathbf{P}.$$

To see this, notice that

$$\pi_{1j} = \sum_{i=1}^n \pi_{0i} \mathbf{P}_{ij}.$$

Iterating, we get the evolution of the unconditional distribution

$$\pi'_t = \pi'_{t-1} \mathbf{P} = \dots = \pi'_0 \mathbf{P}^t. \quad (1.1)$$

Definition 1.2. *An unconditional distribution is called **stationary** if $\pi_{t+1} = \pi_t = \pi$. We say that (\mathbf{P}, π) is a **stationary Markov chain** if the initial distribution is stationary, $\pi_0 = \pi$.*

It follows from (1.1) that a stationary distribution satisfies

$$\pi' = \pi' \mathbf{P} \quad \implies \quad \pi = \mathbf{P}' \pi$$

i.e., π is the left eigenvector of \mathbf{P} (or the (right) eigenvector of \mathbf{P}') associated with eigenvalue equal to 1,

$$(I - \mathbf{P}') \pi = 0.$$

We want to study conditions for the existence and uniqueness of a stationary distribution. Furthermore, we are interested in investigating the limits (if they exist) of the evolution of unconditional distributions

$$\lim_{t \rightarrow \infty} \pi_t \doteq \pi_\infty.$$

Observe that the limit, if it exists, is also a stationary distribution, $\pi'_\infty = \pi'_\infty \mathbf{P}$. Clearly, if the same limit is reached from any initial π_0 , then the stationary distribution is also unique.

Definition 1.3. Let π_∞ be a unique vector that satisfies $(I - \mathbf{P}') \pi_\infty = 0$. If, for all initial distributions π_0 , we have $\lim_{t \rightarrow \infty} (\mathbf{P}^t)' \pi_0 = \pi_\infty$, we say that the Markov chain is **asymptotically stationary with a unique stationary distribution**.

Trivially, if we initiate the Markov chain at the stationary distribution, $\pi_0 = \pi_\infty$, then the Markov chain (\mathbf{P}, π_∞) is stationary in the sense of Definition A.11.

Example 1.1. A Markov chain with the transition matrix given by

$$\mathbf{P} = \begin{bmatrix} 0.7 & 0.3 & 0 \\ 0 & 0.5 & 0.5 \\ 0 & 0.9 & 0.1 \end{bmatrix}.$$

has a unique stationary distribution $\pi' = [0, 0.6429, .3571]$ that can be reached as the asymptotic distribution from any initial distribution, $\lim_{t \rightarrow \infty} (\mathbf{P}^t)' \pi_0 = \pi$. The Markov chain (\mathbf{P}, π_0) is therefore asymptotically stationary with a unique stationary distribution. The set (e_2, e_3) is called the **absorbing subset**—once a path reaches this subset, it can never leave it.

Example 1.2. While a stationary distribution of a Markov chain always exists, it may not be unique. Consider a Markov chain with transition matrix

$$\mathbf{P} = \begin{bmatrix} 1 & 0 & 0 \\ .2 & .5 & .3 \\ 0 & 0 & 1 \end{bmatrix}.$$

There are two **absorbing states** e_1 and e_3 , associated with stationary distributions $(\pi^1)' = [1, 0, 0]$ and $(\pi^3)' = [0, 0, 1]$. Since $\mathbf{P}' \pi^j = \pi^j$, $j = 1, 3$, then also $\mathbf{P}' \tilde{\pi} = \tilde{\pi}$ for

$$\tilde{\pi} = \lambda \pi^1 + (1 - \lambda) \pi^3, \quad \lambda \in [0, 1],$$

so the Markov chain has infinitely many stationary distributions. One of these stationary distributions is always reached as the asymptotic distribution but the particular limit depends on the initial distribution π_0 .

Example 1.3. Finally, consider a Markov chain with the transition matrix

$$\mathbf{P} = \begin{bmatrix} 0 & 0 & 1 \\ .2 & .5 & .3 \\ 1 & 0 & 0 \end{bmatrix}.$$

This Markov chain has a unique stationary distribution $\pi' = [0.5, 0, 0.5]$ but this cannot be asymptotically reached from any initial distribution other than $\pi_0 = \pi$.

Proposition 1.2. Let there be an $m \geq 1$ such that $(\mathbf{P}^m)_{ij} > 0, \forall i, j \in \{1, \dots, n\}$. Then \mathbf{P} has a unique stationary distribution with strictly positive elements and is asymptotically stationary.

The proposition states that if there is enough ‘mixing’ (i.e., the ability to move from each state to each state with a strictly positive probability in finite time, a property called **irreducibility**), the process has a unique stationary distribution that is achieved asymptotically from any initial distribution π_0 .

Proof. The result in the proposition is a consequence of the Perron–Frobenius theorem. First observe that for any eigenvalue λ of \mathbf{P} associated with an eigenvector v , we have $\mathbf{P}^m v = \lambda^m v$, and so λ^m is the eigenvalue of \mathbf{P}^m . Next, since \mathbf{P}^m is a transition matrix, we have $\mathbf{P}^m \mathbf{1} = \mathbf{1}$ where $\mathbf{1} = (1, \dots, 1)'$, and therefore $r = 1$ is an eigenvalue associated with the eigenvector $\mathbf{1}$.

The matrix \mathbf{P}^m has strictly positive elements and $r = 1$ is associated with a strictly positive eigenvector, so then, by the Perron–Frobenius theorem, $r = 1$ is also the unique largest eigenvalue of \mathbf{P}^m and the associated eigenvector is unique. This means that for any other eigenvalue λ^m different from r , we have $|\lambda^m| < 1$, and hence also $|\lambda| < 1$. Therefore, $r = 1$ is also the unique largest eigenvalue of \mathbf{P} (even though \mathbf{P} may not have all strictly positive elements).

Since the eigenvalues and their geometric multiplicity (the number of associated linearly independent eigenvectors, see the second result from Theorem 1.1) of a matrix and its transpose coincide, $r = 1$ is also the unique largest eigenvalue of \mathbf{P}' , with a unique eigenvector with strictly positive elements that sum up to one. This eigenvector yields the unique stationary distribution.

The convergence to the stationary distribution follows from the fact that all eigenvalues of \mathbf{P}' other than $r = 1$ have magnitude smaller than one. To simplify the following arguments, assume that \mathbf{P}' has n linearly independent eigenvectors v_1, \dots, v_n .¹ The eigenvectors then form the so-called **eigenvector base** of the space of all vectors in \mathbb{R}^n , and we can therefore find projection coefficients c^1, \dots, c^n such that

$$\pi_0 = c^1 v_1 + \dots + c^n v_n.$$

Without loss of generality, assume that v_1, \dots, v_n are associated with eigenvalues in descending order, i.e., v_1 is associated with the Perron root $\lambda_1 = r = 1$ and $|\lambda_j| < 1$ for

¹This is true, for example, when \mathbf{P}' has n distinct eigenvalues. The logic of all following arguments goes through even in the case when n linearly independent eigenvectors do not exist, using generalized eigenvectors instead.

$j = 2, \dots, n$. Then

$$\begin{aligned} \mathbf{P}'\pi_0 &= \mathbf{P}'(c^1v_1 + \dots + c^n v_n) = c^1\mathbf{P}'v_1 + \dots + c^n\mathbf{P}'v_n \\ &= c^1\lambda_1v_1 + \dots + c^n\lambda_nv_n \end{aligned}$$

and by iteration

$$\pi_t = (\mathbf{P}')^t \pi_0 = c^1v_1 + \dots + c^n (\lambda_n)^t v_n \rightarrow c^1v_1 \quad (1.2)$$

where, by construction, v_1 has entries that sum to one and are strictly positive, because it is an eigenvector of the strictly positive matrix $(\mathbf{P}')^m$. Only the eigenvector associated with $r = 1$ survives iteration on \mathbf{P} , and therefore constitutes the stationary distribution. Also notice that the stationary distribution assigns strictly positive probabilities to all states of the Markov chain. ■

Equation (1.2) shows that, in fact, the projection component $c^1v_1 = \pi$. Necessarily, $c^1 > 0$, because the remaining components vanish to zero, and then π_t could not be a probability distribution. The remaining components $(\lambda_i)^t c^i v_i$, $i \geq 2$ then represent transitory deviations from the stationary distribution that vanish at geometric rates λ_i . In fact, each of the vectors v_i , $i \geq 2$ must have elements that sum up to zero, so that π_t remains a probability distribution as t increases.

Returning back to Examples 1.1 to 1.3, neither of them satisfies the assumption of Proposition 1.2 that there exists an $m \geq 1$ such that \mathbf{P}^m has all strictly positive elements. This has different implications in the three examples.

In Example 1.1, the eigenvalue $r = 1$ has a unique associated eigenvector and all other eigenvalues satisfy $|\lambda| < 1$, so there still exists a unique stationary distribution that can be reached asymptotically from any initial π_0 but the stationary distribution fails to have a full support.

In Example 1.2, the eigenvalue $r = 1$ has two associated strictly positive eigenvectors, so we get a multiplicity of stationary distributions.

In Example 1.3, the eigenvalue $r = 1$ does have a unique eigenvector so we get a unique stationary distribution but there is another eigenvalue $\lambda = -1$, so the argument in equation (1.2) fails and convergence to the stationary distribution is not guaranteed. To see this, consider again equation (1.2):

$$\pi_t = (\mathbf{P}')^t \pi_0 = c^1 (\lambda_1)^t v_1 + c^2 (\lambda_2)^t v_2 + \dots + c^n (\lambda_n)^t v_n$$

Now assume that $\lambda_1 = 1$, $|\lambda_2| = 1$ but $\lambda_2 \neq 1$, and $|\lambda_j| < 1$ for $j \geq 3$. Then, as $t \rightarrow \infty$,

$$\pi_t \approx c^1v_1 + c^2 (\lambda_2)^t v_2.$$

Since $(\lambda_2)^t$ does not converge, then π_t also does not converge, unless π_0 is chosen in a way that the projection coefficient $c^2 = 0$. When $c^2 \neq 0$, π_t will asymptotically contain a cycle.

1.1.3 Forecasting and conditional expectations

We are interested in computing conditional expectations of random variables. Start with

$$E [x_{t+1} | x_t = e_i] = \sum_{j=1}^n e_j P(x_{t+1} = e_j | x_t = e_i) = \sum_{j=1}^n e_j \mathbf{P}_{ij} = (\mathbf{P}_{i \cdot})' = \mathbf{P}' e_i.$$

Here, the state x_{t+1} realized in period $t + 1$ is random from the perspective of time t , and the conditional expectation yields the vector of conditional probabilities of the state tomorrow, conditional on $x_t = e_i$. Denote \bar{y} an $n \times 1$ vector of numbers. Then $y_t = \bar{y}' x_t$ is a random variable, with realization \bar{y}_j if state $x_t = e_j$ is realized. We then have

$$E [y_{t+1} | x_t = e_i] = \bar{y}' E [x_{t+1} | x_t = e_i] = \bar{y}' (\mathbf{P}_{i \cdot}) = (\mathbf{P} \bar{y})_i = \bar{y}' \mathbf{P}' e_i = e_i' \mathbf{P} \bar{y}$$

Similarly

$$E [y_{t+k} | x_t = e_i] = E [\bar{y}' x_{t+k} | x_t = e_i] = (\mathbf{P}^k)_i \cdot \bar{y} = (\mathbf{P}^k \bar{y})_i = e_i' \mathbf{P}^k \bar{y} \quad (1.3)$$

Hence, for a general state x_t , we can write

$$E [y_{t+k} | x_t] = x_t' \mathbf{P}^k \bar{y}$$

Remark 1.1. *There is a simple heuristic for the left and right multiplication of \mathbf{P} . Left multiplication implies rolling the Markov chain forward, while right multiplication implies rolling the chain backward.*

For instance, starting with a probability distribution π this period, $\pi' \mathbf{P}$ is the probability distribution next period. On the other hand, given the vector \bar{y} of realizations of a random variable next period, $\mathbf{P} \bar{y}$ is the conditional expectation.

Hence the term $x_t' \mathbf{P}^k \bar{y}$ can be interpreted as follows. $x_t \mathbf{P}^k$ is the conditional distribution of the Markov chain at $t + k$, given state t today. Then $x_t \mathbf{P}^k \bar{y}$ is the conditional expectation of a random variable $y_{t+k} = \bar{y}' x_{t+k}$, conditional on x_t today.

In what follows, we will also be interested in the conditional expectation of a random variable conditional on x_t being in a particular set of states $\mathcal{X} = \{e_{i_1}, \dots, e_{i_k}\}$. Following Definition A.9, we construct the probability vector $\pi_t^{\mathcal{X}}$ with elements

$$\pi_{t,i}^{\mathcal{X}} = \frac{\pi_{t,i} \mathbf{1}\{e_i \in \mathcal{X}\}}{\sum_{j: e_j \in \mathcal{X}} \pi_{t,j}} \quad (1.4)$$

as the unconditional probability distribution of the Markov chain at time t , restricted to the set \mathcal{X} and normalized to assure that $\pi_t^{\mathcal{X}}$ is a probability vector. Then

$$E [y_{t+k} | x_t \in \mathcal{X}] = (\pi_t^{\mathcal{X}})' \mathbf{P}^k \bar{y}.$$

Two special cases are $\mathcal{X} = \{e_i\}$, in which case we recover conditional expectation (1.3),

and $\mathcal{X} = \{e_1, \dots, e_n\}$, which yields the unconditional expectation

$$E[y_{t+k}] = \pi_t' \mathbf{P}^k \bar{y} = \pi_{t+k}' \bar{y}.$$

1.1.4 Invariant functions and ergodicity

In time series econometrics, we would like to associate unconditional distributions with time-series averages. We therefore want to establish conditions under which the stationary distribution is unique and the mean of a random variable under this stationary distribution corresponds to the time-series average of a long sample of realizations of this random variable on the Markov chain. This will lead to a so called **ergodic theorem**, which is a generalization of the law of large numbers for stochastic processes.

The following example illustrates a situation that constitutes a problem for empirical work.

Example 1.4. Consider the Markov chain given by

$$\mathbf{P} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \pi_0 = \pi = \begin{bmatrix} p & 1-p \end{bmatrix}, \quad p \in [0, 1].$$

Every path of the random variable $y_t = \bar{y}' x_t$ is constant, with realization \bar{y}_i that depends on the initial draw of the Markov chain. The cross-sectional distribution of y_t under the stationary distribution has realizations \bar{y}_1 and \bar{y}_2 , with distribution π . When $\bar{y}_1 \neq \bar{y}_2$ and $p \in (0, 1)$, then the time series average along a single observed infinite path, equal to

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T y_t = \bar{y}_i,$$

differs from the cross-sectional prediction $E[y_t] = \pi' \bar{y}$.

We define the concept of **ergodicity** using so-called invariant functions.

Definition 1.4. Let (\mathbf{P}, π) be a stationary Markov chain. A random variable $y_t = \bar{y}' x_t$ is said to be **invariant** if $y_t = y_0$, $t \geq 0$ for all realizations of x_t , $t \geq 0$ that occur with positive probability under (\mathbf{P}, π) .

A random variable is therefore invariant if it remains constant along a set of paths drawn from the Markov chain that has probability one. However, that does not mean that \bar{y} must be a constant vector because y_0 may differ depending on the initiation of x_0 . We will equivalently refer to random variables $y_t = \bar{y}' x_t$ satisfying Definition 1.4 as invariant random variables, and to the associated vectors \bar{y} as invariant functions. For example, in Example 1.4, any function $\bar{y} = (\bar{y}_1, \bar{y}_2)'$ is invariant.

The following proposition illustrates the role of invariant functions.

Proposition 1.3. Let (\mathbf{P}, π) be a stationary Markov chain. If

$$E[y_{t+1} | x_t] = y_t \tag{1.5}$$

then the random variable $y_t = \bar{y}'x_t$ is invariant.

Proof. The Markov chain (\mathbf{P}, π) implies an unconditional distribution over two subsequent states (x_{t+1}, x_t) . Since y_t is a deterministic function of x_t , applying the Law of Iterated Expectations yields

$$\begin{aligned} E \left[(y_{t+1} - y_t)^2 \right] &= E \left[E \left[y_{t+1}^2 - 2y_t y_{t+1} + y_t^2 \mid x_t \right] \right] \\ &= E \left[E \left[y_{t+1}^2 \mid x_t \right] - 2y_t E \left[y_{t+1} \mid x_t \right] + y_t^2 \right] \\ &= E \left[y_{t+1}^2 \right] - 2E \left[y_t^2 \right] + E \left[y_t^2 \right] = 0. \end{aligned}$$

Hence the set of paths for which $y_{t+1} = y_t$ has probability one. ■

A random variable satisfying (1.5) is called a **martingale**. Martingales are processes that neither grow or decay on average, conditional on the state today. The statement of the above proposition says that a martingale on a stationary Markov chain must have constant paths. This is closely related to the so-called **martingale convergence theorem**.

Observe that (1.5) can also be written as

$$x_t' \mathbf{P} \bar{y} = x_t' \bar{y} \quad \forall x_t$$

and hence

$$\mathbf{P} \bar{y} = \bar{y}$$

Therefore, \bar{y} must be the right eigenvector of \mathbf{P} associated with a unit eigenvalue. One trivial solution is $\bar{y} = \mathbf{1}$ but we are interested in what happens when there are multiple linearly independent solutions.

Definition 1.5. Let (\mathbf{P}, π) be a stationary Markov chain. The chain is said to be **ergodic** if the only invariant functions \bar{y} are constant with probability 1 under the stationary probability distribution π , i.e., $\bar{y}_i = \bar{y}_j$ for all i, j with $\pi_i, \pi_j > 0$.

Definition 1.6. A set \mathcal{X} of states of a stationary Markov chain (\mathbf{P}, π) is called **invariant** if it cannot be entered from any state $e_i \notin \mathcal{X}$ that has a strictly positive probability under π :

$$\{e_i : P(x_{t+1} \in \mathcal{X} \mid x_t = e_i) \pi(e_i) > 0\} \subseteq \mathcal{X}.$$

Since an invariant set of a Markov chain cannot be entered, then, if it has a strictly positive measure under the stationary density π , it cannot be exited with a positive probability, either. Invariant sets that are disjoint and have strictly positive measure under π therefore divide the state space into components that do not communicate with each other. If such disjoint invariant sets with strictly positive measure exist, then the Markov chain is not ergodic.

Lemma 1.4. A stationary Markov chain (\mathbf{P}, π) is ergodic if and only if it has no disjoint invariant sets that have a strictly positive measure.

Proof. Let there be two disjoint invariant sets \mathcal{X}_1 and \mathcal{X}_2 of strictly positive measure under π . Define $y_t = \bar{y}'x_t$ with $\bar{y}_i = \mathbf{1}\{e_i \in \mathcal{X}_1\}$. Then \bar{y} is an invariant function, namely if

$x_0 \in \mathcal{X}_1$, then $y_t = 1$, while if $x_0 \notin \mathcal{X}_1$, then $y_t = 0$. We have an invariant function that is not constant under the stationary distribution π , and (\mathbf{P}, π) is not ergodic.

On the other hand, assume that (\mathbf{P}, π) is not ergodic, and take the associated invariant function \bar{y} that is not constant under π . Denote $\{\hat{y}_k\}$ the set of distinct values of elements of the vector \bar{y} and construct sets $\mathcal{X}_k = \{e_i : \bar{y}_i = \hat{y}_k, i = 1, \dots, n\}$. Since \bar{y} is not constant under π , at least two such sets \mathcal{X}_k have strictly positive measure under π . Moreover, these two sets are invariant. To see this, assume it is possible to enter \mathcal{X}_k from another \mathcal{X}_l that has a strictly positive probability. Then it implies a path y_t that is not constant, a contradiction to \bar{y} being an invariant function. Hence we have found two disjoint invariant sets of strictly positive measure under π . ■

In order to verify ergodicity, it therefore suffices to characterize the invariant sets of the Markov chain. With these results at hand, we can state the **ergodic theorem** for a stationary Markov chain.

Proposition 1.5 (Birkhoff–Khinchin). *Let (\mathbf{P}, π) be a stationary Markov chain and $y_t = \bar{y}'x_t$. Fix an initial state x_0 and let \mathcal{X} be the smallest invariant set to which x_0 belongs, with $\pi^\mathcal{X}$ being the conditional probability distribution defined in (1.4). Then*

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T y_t = \left(\pi^\mathcal{X} \right)' \bar{y} \quad (1.6)$$

with probability one. Moreover, if the Markov chain is also ergodic, then the right-hand side is equal to $\pi' \bar{y}$.

Proof. For a discussion and proof in a more general setting, for example, Billingsley (1979), Chapter 24. ■

Proposition 1.5 states that in general, sample averages along paths may depend on initial conditions. However, if we partition the state space into its smallest invariant set, then paths initiated within each set share the same long-run statistical properties and hence sample averages.

When the Markov chain is ergodic, then the smallest invariant set \mathcal{X} to which x_0 belongs has measure one under π , and hence $\pi^\mathcal{X} = \pi$. Then the time-series average along every path converges to the unconditional expectation of y_t under the stationary distribution. This is a manifestation of the so-called **mixing condition**, which states that dependence on initial conditions vanishes if the states of the Markov chain communicate with each other.

In Example 1.4, the Markov chain is not ergodic when $p \in (0, 1)$, and there are two non-trivial invariant sets $\mathcal{X} = \{e_1\}$ and $\mathcal{X} = \{e_2\}$. For example, when $x_0 = e_1$, then $\pi_t^\mathcal{X}$ in equation (1.6) is equal to $[1, 0]'$, and the right-hand side is equal to \bar{y}_1 , and analogously for $x_0 = e_2$. Observe that the Markov chain from Example 1.4 satisfies the ergodicity definition when $p \in \{0, 1\}$.

Corollary 1.6. *Let (\mathbf{P}, π_0) be an asymptotically stationary Markov chain with a unique stationary*

distribution π . Then

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T y_t = E \left[\lim_{t \rightarrow \infty} y_t \mid x_0 \right] = \lim_{t \rightarrow \infty} E[y_t], \quad (1.7)$$

where the first equality holds with probability one, and the second limit represents convergence in distribution. Further, the stationary Markov chain (\mathbf{P}, π) is ergodic.

Proof. Since the unique stationary distribution π is achieved from any initial condition, the smallest invariant set has probability one and the Markov chain (\mathbf{P}, π) is ergodic. The distribution of x_t converges to the unique π , so the limiting distribution of y_t uniquely exists, and hence

$$E \left[\lim_{t \rightarrow \infty} y_t \mid x_0 \right] = \lim_{t \rightarrow \infty} (\pi_t)' \bar{y} = \pi' \bar{y}.$$

■

Ergodicity is a property of the pair (\mathbf{P}, π) , where π is a stationary distribution. When (\mathbf{P}, π) is ergodic, it does not mean that π is the unique stationary distribution associated with \mathbf{P} . Example 1.2 for $\lambda \in \{0, 1\}$ is such a case. However, if \mathbf{P} has a unique stationary distribution π , then (\mathbf{P}, π) is ergodic. Equivalently, take (\mathbf{P}, π) which is not ergodic, then π is not a unique stationary distribution. This follows from Lemma 1.4. A stationary Markov chain (\mathbf{P}, π) that is not ergodic has two invariant sets \mathcal{X}_1 and \mathcal{X}_2 of a strictly positive measure under π , and hence conditional distributions $\pi^{\mathcal{X}_1}$ and $\pi^{\mathcal{X}_2}$ defined as restrictions of π using (1.4) are also stationary.

Example 1.5. Consider a Markov chain with transition matrix

$$\mathbf{P} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

This Markov chain has a unique stationary distribution $\pi' = [0.5, 0.5]$. Then (\mathbf{P}, π) is ergodic, with the smallest non-empty invariant set $\mathcal{X} = \{e_1, e_2\}$. Realizations of y_t along paths cycle deterministically between \bar{y}_1 and \bar{y}_2 and the limit of the time series average is equal to $\frac{1}{2}\bar{y}_1 + \frac{1}{2}\bar{y}_2$, which is equal to $(\pi^{\mathcal{X}})' \bar{y}$, in line with the ergodic theorem.

However, the Markov chain is not asymptotically stationary with a unique stationary distribution, so Corollary 1.6 does not apply. Indeed, the limits $\lim_{t \rightarrow \infty} y_t$ and $\lim_{t \rightarrow \infty} E[y_t]$ in equation (1.7) do not exist in this case when $\bar{y}_1 \neq \bar{y}_2$.

Ergodicity and asymptotic behavior of Markov chains is closely related to the characterization of eigenvalues associated with the transition matrix \mathbf{P} , i.e., solutions to the eigenvalue equation

$$\mathbf{P}v = \lambda v$$

that lie on the unit circle, $|\lambda| = 1$.

Since \mathbf{P} is a transition matrix, it has an eigenvalue equal to one, simply because $\mathbf{P}\mathbf{1} = \mathbf{1}$. Eigenvalues of a matrix and its transpose coincide, so there must exist a nonzero solution

π with non-negative elements to the equation

$$\mathbf{P}'\pi = \pi. \quad (1.8)$$

The fact that the solution is non-negative is the consequence of an extension of the Perron–Frobenius theorem to non-negative matrices. Hence there must exist a stationary distribution for the Markov chain. Another consequence of the extension of the Perron–Frobenius theorem is that there are no eigenvalues of \mathbf{P} that are larger than one in magnitude but there are potentially some with magnitude equal to one.

To illustrate the implications of these results, we return to Examples 1.1–1.3.

The first and simplest case is the situation when the solution to (1.8) is unique (up to a normalization), and there are no other eigenvalues on the unit circle. This is the case from Example 1.1. In this case, the Markov chain (\mathbf{P}, π_0) is asymptotically stationary with a unique stationary distribution π , as shown in the proof of Proposition 1.2. By Corollary 1.6, the stationary Markov chain (\mathbf{P}, π) is then ergodic. Indeed, the invariant set $\mathcal{X} = \{e_2, e_3\}$ has measure one under π and cannot be further partitioned into invariant subsets.

A sufficient condition for this case is when \mathbf{P} has all strictly positive elements (or, more generally, there exists an m such that \mathbf{P}^m has all strictly positive elements), as assumed in Propositions 1.1 and 1.2. When this sufficient condition holds, then, in addition, the stationary distribution will have strictly positive mass in all states.

The second case occurs when there are multiple linearly independent solutions to (1.8), i.e., when the unit eigenvalue has geometric multiplicity higher than one. This is the case from Example 1.2. Then there exist multiple (in fact, infinitely many) stationary distributions π , and, if there are no other eigenvalues on the unit circle, one of them will be reached as the limit $\lim_{t \rightarrow \infty} \pi_0' \mathbf{P}^t$. This limit generally depends on π_0 . Whether the stationary Markov chain (\mathbf{P}, π) is ergodic may differ across the alternative limiting distributions π . For example, in Example 1.2, the Markov chain is ergodic for $\lambda \in \{0, 1\}$ and is not ergodic for $\lambda \in (0, 1)$. In the latter case, we have invariant sets $\mathcal{X}_1 = \{e_1\}$ and $\mathcal{X}_2 = \{e_3\}$ with strictly positive probability.

Finally, the third case is a situation when there is also another eigenvalue on the unit circle that is different from one. Then the unconditional distribution π_t may not generally converge to a stationary distribution. This is the case from Example 1.3, in which there exists a unique stationary distribution π and the stationary chain (\mathbf{P}, π) , despite the cyclical properties of the paths of x_t , is ergodic because the invariant set $\mathcal{X} = \{e_1, e_3\}$ has measure one under π and cannot be further partitioned into invariant subsets. However, Corollary 1.6 does not apply.

1.1.5 Simulation and likelihood estimation

It is easy to construct simulated paths from a Markov chain. If the current draw is $x_t = e_i$, then the row \mathbf{P}_i represents the conditional probability of x_{t+1} . Then we need to draw a random variable uniformly distributed on $[0, 1]$ and assign x_{t+1} according to which bracket of the cdf constructed from \mathbf{P}_i the random draw falls into.

We may now ask the opposite question. If we observe a particular sample $\{\hat{x}_t\}_{t=0}^T$, what is the probability that this sample has been drawn from a particular Markov chain (\mathbf{P}, π_0) ? Let the particular sample of observed data be

$$\hat{x}^T = (\hat{x}_0, \hat{x}_1, \dots, \hat{x}_T) = (e_{i_0}, e_{i_1}, \dots, e_{i_T}).$$

and let us assume that there is a family of candidate Markov chains on the table, indexed by a parameter vector $\theta \in \Theta$ such that $\pi_0 = \pi_0(\theta)$ and $\mathbf{P} = \mathbf{P}(\theta)$. Each pair $(\mathbf{P}(\theta), \pi_0(\theta))$ thus defines a probability measure $P(\theta)$.

The likelihood of the data \hat{x}^T for a given parameter vector is

$$L(\theta|\hat{x}^T) = P(\hat{x}_0, \hat{x}_1, \dots, \hat{x}_T|\theta) = \pi_{0,i_0}(\theta) \mathbf{P}_{i_0,i_1}(\theta) \mathbf{P}_{i_1,i_2}(\theta) \dots \mathbf{P}_{i_{T-1},i_T}(\theta).$$

Denote n_{ij} the number of transitions from state e_i to e_j . Then the likelihood has a multinomial distribution

$$L(\theta|\hat{x}^T) = \pi_{0,i_0}(\theta) \prod_{i,j=1}^n (\mathbf{P}_{i,j}(\theta))^{n_{ij}}.$$

The **inference problem** of the econometrician is to deduce, from the observed sample $\{\hat{x}_t\}_{t=0}^T$, the probability measure $P(\theta_0)$ that generated the data, where $\theta_0 \in \Theta$ is the 'true' parameter value.

One standard method is based on maximizing the likelihood function. The **maximum likelihood estimator** corresponds to maximizing $L(\theta|\hat{x}^T)$ with respect to θ :

$$\hat{\theta}^{ML} = \arg \max_{\theta \in \Theta} L(\theta|\hat{x}^T).$$

Often, given the multiplicative structure of the likelihood, it is more sensible to maximize

$$\log L(\theta|\hat{x}^T) = \log \pi_{0,i_0}(\theta) + \sum_{i,j=1}^n n_{ij} \log \mathbf{P}_{i,j}(\theta). \quad (1.9)$$

The maximum likelihood estimator has a range of desirable statistical properties because it is based on the maximization of the likelihood function, which summarizes all information about the underlying probability distribution contained in the data. On the other hand, it requires a complete specification of this probability distribution and it can be hard to compute in more complex models. Moreover, the econometrician may be concerned that some parts of the model may be misspecified or trust certain distributional assumptions less than others, and therefore would like to omit such information from the estimation.

1.1.6 Method of moments estimation

The method of moments is based on the idea that we can construct a vector of conditions, called moments, that the distribution of the stochastic process satisfies only at the true parameter value θ_0 . Solving the empirical counterpart of these moment conditions yields

the **method of moments estimator** $\hat{\theta}^{MOM}$.

Specifically, let there be a $K \times 1$ vector $g(x; \theta) = (g_1(x; \theta), \dots, g_K(x; \theta))'$ of functions that satisfy

$$E[g(x_t; \theta)] = 0$$

if and only if $\theta = \theta_0$. To implement the estimator, we replace the theoretical expectation under the stationary distribution with the sample average and solve the vector of **moment conditions** for $\hat{\theta}^{MOM}$:

$$\frac{1}{T+1} \sum_{t=0}^T g(\hat{x}_t; \hat{\theta}^{MOM}) = 0.$$

There is generally no unique way how to pick the vector of moments. In practice, the choice is motivated by the tradeoff between statistical efficiency, complexity of the construction of the moment condition, and judgement of the econometrician concerning the correct specification of different parts of the model. The **generalized method of moments** devised by Hansen (1982) provides a formal way how to combine moment conditions efficiently and how to assess the statistical properties of the moments estimator.

In Section 4.3.1, we revisit the estimation problem and contrast the above two methods with a Bayesian approach to estimation.

Example 1.6. Consider a two-state Markov chain with the transition matrix

$$\mathbf{P} = \begin{bmatrix} 1-p & p \\ q & 1-q \end{bmatrix}$$

with $p > 0$, $q > 0$. Assume that q is known, and we are estimating the parameter $\theta = p$ from a sample of data \hat{x}^T , which was generated from a Markov chain with the true parameter value p_0 , hence $\theta_0 = p_0$.

Using (1.9), the log-likelihood function for the data is given by

$$\log L(p|\hat{x}^T) = \log \pi_{0,i_0} + n_{11} \log(1-p) + n_{12} \log p + n_{21} \log q + n_{22} \log(1-q).$$

where n_{ij} counts the number of transitions from state e_i to state e_j in the sample \hat{x}^T . The first-order condition with respect to p yields

$$0 = -n_{11} \frac{1}{1-p} + n_{12} \frac{1}{p},$$

and hence the **maximum likelihood estimator** is given by

$$\hat{p}^{ML} = \frac{n_{12}}{n_{11} + n_{12}}.$$

Since the estimator \hat{p}^{ML} computes the share of cases when the chain that is currently in state e_1 transited into state e_2 , the Law of Large Numbers implies that $\hat{p}^{ML} \rightarrow p_0$.

One example of a **method of moments estimator** can be constructed as follows. The stationary

distribution of the Markov chain is given by

$$\pi' = \left(\frac{q}{p+q}, \frac{p}{p+q} \right).$$

Since the Markov chain is ergodic, we can rely on the result from the ergodic theorem, Theorem 1.6. Construct the moment function

$$g(x; p) = \mathbf{1}\{x = e_1\} - \frac{q}{p+q}.$$

Observe that

$$E[g(x; p)] = \pi_1 - \frac{q}{p+q} = 0,$$

and, if the moment π_1 was generated using the true parameter, then there is indeed a unique value $p = p_0$ that solves this moment condition. The empirical counterpart to the moment condition is

$$\frac{1}{T+1} \sum_{t=0}^T \mathbf{1}\{\hat{x}_t = e_1\} - \frac{q}{p+q} \doteq \hat{\pi}_1 - \frac{q}{p+q} = 0.$$

By the ergodic theorem, $\hat{\pi}_1 \rightarrow \pi_1$ as $T \rightarrow \infty$, and we can therefore construct the method of moments estimator as

$$\hat{p}^{\text{MOM}} = \frac{1 - \hat{\pi}_1}{\hat{\pi}_1} q.$$

1.1.7 Hidden Markov models

We can use Markov chains to derive a version of a frequently used macroeconomic framework, so-called **Hidden Markov models**.

Recall that

$$E[x_{t+1} | x_t] = \mathbf{P}'x_t.$$

Then we can derive the mean-zero residual

$$v_{t+1} \doteq x_{t+1} - E[x_{t+1} | x_t] = x_{t+1} - \mathbf{P}'x_t.$$

Hence we get a 'law of motion' for the state x_t

$$x_{t+1} = \mathbf{P}'x_t + v_{t+1}.$$

Let us now assume that x_t is 'hidden', i.e., not observable by the econometrician. Instead of x_t , the econometrician observes a noisy signal about x_t . Specifically, let y_t be a random realization of a coordinate vector f_j from an m -dimensional space.

The observation matrix is

$$P(y_t = f_j | x_t = e_i) = \mathbf{Q}_{ij}$$

where \mathbf{Q} is an $n \times m$ matrix, with row entries summing up to one. When \mathbf{Q} has multiple strictly positive entries in a given column, it means that the same signal realization can occur in multiple different states, leading to information loss. This can happen when $m < n$, and we have less signal realizations than states, or when the state is observed with noise, and we have multiple strictly positive entries in a given row of \mathbf{Q} .

We can calculate the expected value of the signal conditional on the true state as

$$E[y_t | x_t = e_i] = (\mathbf{Q}_i)',$$

and therefore

$$E[y_t | x_t] = \mathbf{Q}'x_t.$$

Again, we can define the mean-zero residual

$$u_t = y_t - E[y_t | x_t].$$

Combining these equations, we obtain the system

$$\begin{aligned} x_{t+1} &= P'x_t + v_{t+1} \\ y_t &= Q'x_t + u_t. \end{aligned} \tag{1.10}$$

The first equation is the **law of motion** for the unobservable state x_t . The second equation is a '**measurement equation**' for observed data y_t . Our goal is to use the observed data to figure out what the evolution of the hidden state is. Observe that due to the noise u_t , this inference will be imperfect. We will be solving a **filtering problem**, i.e., we will use observations y_t and attempt to separate noise u_t from the best forecast of the evolution of the state x_t .

Nonlinear filtering

Imagine we observed the sequence $y^t = (y_1, \dots, y_t)$. What is the best forecast of x_t , given this sequence? Formally, denote

$$\rho_i(t) = P(x_t = e^i | y^t) = P(x_t = e^i | y_1, y_2, \dots, y_t).$$

Stacking $\rho_i(t)$, $i = 1, \dots, n$ in a column vector $\rho(t)$, we obtain

$$\rho(t) = P(x_t | y^t).$$

We want to find a recursive representation for $\rho(t)$. In particular, we want to show that we can express $\rho(t)$ by combining:

- $\rho(t-1) = P(x_{t-1} | y^{t-1})$, i.e., the best forecast of x_{t-1} given all data y^{t-1} up to time t ,
- the law of motion for the state, $x_t = P'x_{t-1} + v_t$, consisting of the best forecast of x_t

given x_{t-1} , and the associated error v_t ,

- and incorporating the new observation y_t whose conditional probability is $P(y_t | x_t)$.

Let

$$\begin{aligned}
 P(x_t | y^t) &= \frac{P(x_t, y_t | y^{t-1})}{P(y_t | y^{t-1})} = \frac{\sum_{x_{t-1}} P(x_t, x_{t-1}, y_t | y^{t-1})}{\sum_{x_t} \sum_{x_{t-1}} P(x_t, x_{t-1}, y_t | y^{t-1})} \\
 &= \frac{\sum_{x_{t-1}} P(y_t | x_t, x_{t-1}, y^{t-1}) P(x_t, x_{t-1} | y^{t-1})}{\sum_{x_t} \sum_{x_{t-1}} P(y_t | x_t, x_{t-1}, y^{t-1}) P(x_t, x_{t-1} | y^{t-1})} \\
 &= \frac{\sum_{x_{t-1}} P(y_t | x_t) P(x_t | x_{t-1}, y^{t-1}) P(x_{t-1} | y^{t-1})}{\sum_{x_t} \sum_{x_{t-1}} P(y_t | x_t) P(x_t | x_{t-1}, y^{t-1}) P(x_{t-1} | y^{t-1})} \\
 &= \frac{\sum_{x_{t-1}} P(y_t | x_t) P(x_t | x_{t-1}) P(x_{t-1} | y^{t-1})}{\sum_{x_t} \sum_{x_{t-1}} P(y_t | x_t) P(x_t | x_{t-1}) P(x_{t-1} | y^{t-1})}
 \end{aligned}$$

Here, \sum_{x_t} denotes the sum over all realizations e_i of the state vector x_t . Observe that we now have a recursive representation between $P(x_t | y^t)$ and $P(x_{t-1} | y^{t-1})$, given the distribution of observations conditional on state $P(y_t | x_t)$, and distribution of state conditional on the state previous period $P(x_t | x_{t-1})$.

Assume that $x_t = e_i$ and $y_t = f_j$. Then the derivation above can be continued in matrix form as

$$\begin{aligned}
 \rho_i(t) &= P(x_t = e^i | y^t, y_t = f_j) = \frac{\sum_s \mathbf{Q}_{ij} \mathbf{P}_{si} \rho_s(t-1)}{\sum_i \sum_s \mathbf{Q}_{ij} \mathbf{P}_{si} \rho_s(t-1)} = \\
 &= \frac{\sum_s \mathbf{P}_{si} \mathbf{Q}_{ij} \rho_s(t-1)}{\sum_s [\mathbf{PQ}]_{sj} \rho_s(t-1)} = \frac{(\rho(t-1))' \mathbf{P}_{:i} \mathbf{Q}_{ij}}{(\rho(t-1))' [\mathbf{PQ}]_{:j}}
 \end{aligned}$$

where $\rho_s(t-1)$ in the last expression is the column vector with elements $\rho_s(t-1)$. This can be written for an arbitrary $y_t = f_j$ as

$$\rho_i(t) = \frac{(\rho(t-1))' \mathbf{P}_{:i} \mathbf{Q}_{i \cdot} y_t}{(\rho(t-1))' \mathbf{PQ} y_t}$$

and the vector $\rho(t)$, one element for each $x_t = e_i$,

$$\rho(t) = \frac{\text{vec}_i [(\rho(t-1))' \mathbf{P}_{:i} \mathbf{Q}_{i \cdot} y_t]}{(\rho(t-1))' \mathbf{PQ} y_t} = \frac{\text{diag} [\mathbf{Q} y_t (\rho(t-1))' \mathbf{P}]}{(\rho(t-1))' \mathbf{PQ} y_t}$$

where the last line presents two alternative ways of writing the result, where $\text{vec}_i[\cdot]$ is the column vector created by stacking elements indexed by i , while $\text{diag}[\cdot]$ is a column vector created from diagonal elements of the argument. To see that the last expression is correct, notice that the last expression can be written as

$$\rho(t) = \frac{\text{diag} [\mathbf{Q} y_t (\rho(t-1))' \mathbf{P}]}{\text{tr} [(\rho(t-1))' \mathbf{PQ} y_t]} = \frac{\text{diag} [\mathbf{Q} y_t (\rho(t-1))' \mathbf{P}]}{\text{tr} [\mathbf{Q} y_t (\rho(t-1))' \mathbf{P}]} \quad (1.11)$$

so that $\rho(t)$ sum up to one, which is consistent with $\rho(t)$ being a vector of probabilities.

We can now verify the recursivity of the formulation, and how $\rho(t)$ combines

- the previous period best forecast $\rho(t-1)$,
- the law of motion for the state, through the transition matrix \mathbf{P} ,
- the new data point y_t , through the observation matrix \mathbf{Q} .

The Markov property

Notice that while the original system (1.10) was Markov in x_t , the new system cannot because we do not observe x_t . However, we can establish the Markov property for the joint process $\{\rho(t-1), y_t\}$ under a particular filtration.

In particular, let $\rho(0) = P(x_0)$, and denote $\{\tilde{\mathcal{F}}_t\}_{t=0}^{\infty}$ the filtration generated by $\rho(0)$ and the history of random variables y_t , $t = 1, 2, \dots$. Notice that this filtration is coarser (less informative) than the filtration generated by histories of (x_t, y_t) . The claim is that the process $\{\rho(t-1), y_t\}$ is Markov under $\{\tilde{\mathcal{F}}_t\}_{t=0}^{\infty}$, i.e., that the probability distribution of $(\rho(t+j-1), y_{t+j})$ conditional on $\tilde{\mathcal{F}}_t$ is the same as conditional on $(\rho(t-1), y_t)$.

This can be directly seen from the formulas that we derived. First notice that $\rho(t) = P(x_t | y^t)$ is a function of y^t and $\rho(t-1)$ given by equation (1.11). Second, y_{t+1} is a function of x_{t+1} plus random noise, and x_{t+1} itself is a function of x_t plus random noise. Since all information about x_t given data y^t is summarized in $\rho(t)$ and this is a function of y^t and $\rho(t-1)$, then indeed the conditional distribution of y_{t+1} under $\tilde{\mathcal{F}}_t$ is completely characterized by $(\rho(t-1), y_t)$. This establishes the Markov property.

Notice that this conclusion does not mean that the stochastic process $\{\rho(t-1), y_t\}$ is Markov under the original filtration $\{\mathcal{F}_t\}_{t=0}^{\infty}$. This filtration includes information from observations of x_t , and x_t itself is more informative about y_{t+1} than the pair $(\rho(t-1), y_t)$.

Finally, it is important to observe what is the object for which we established the Markov property. It combines the observed data y_t with the evolution of the whole probability distribution $\rho(t) = P(x_t | y^t)$. On the Markov chain, this probability distribution is an n -dimensional vector, and so the problem remains tractable. However, in continuous state spaces, keeping track of the evolution of the whole probability density will not be tractable in most cases. A tractable example will be the case of linear Gaussian Markov models where the filtering problem will reduce to a law of motion of a Markov process representing the first two moments of the distribution. This will be the Kalman filter described in Section 4.1.

1.1.8 Continuous-state Markov chain

The construction of Markov chains can be extended to continuous state spaces. Denote \mathcal{S} a set of states with typical element $s \in \mathcal{S}$. Typically, $\mathcal{S} \subseteq \mathbb{R}^n$. Stochastic evolution of the

state can be described by cumulative distribution functions

$$\begin{aligned}\Pi(s' | s) &= P(s_{t+1} \leq s' | s_t = s), \\ \Pi_0(s) &= P(s_0 \leq s)\end{aligned}$$

The transition density, if it exists, is given by

$$\begin{aligned}\pi(s' | s) &= \frac{d}{ds'} \Pi(s' | s) \\ \pi_0(s) &= \frac{d}{ds} \Pi_0(s)\end{aligned}$$

When the notation is not ambiguous, we also use the transition probability notation

$$\Pi(B | s) = P(s_{t+1} \in B | s)$$

In line with the constructions of the discrete-state Markov chains, we define stationarity, invariant functions and ergodicity.

Definition 1.7. A continuous-state Markov chain (π, π_0) is **stationary** if π_0 satisfies

$$\pi_0(s') = \int \pi(s' | s) \pi_0(s) ds.$$

A function $\phi(s)$ is **invariant** if

$$\int \phi(s') \pi(s' | s) ds' = \phi(s).$$

The Markov chain is **ergodic** if the only invariant functions $\phi(s)$ are constant with probability 1 under a given stationary distribution π_∞ .

Remark 1.2. If the density does not exist then the stationarity condition reads that for every measurable B ,

$$\Pi_0(B) = \int \Pi(s' \in B | s) \Pi_0(ds)$$

where we abused the notation a bit and used Π for the probability measure as opposed to cdf.

Proposition 1.7. Let $y(s)$ be a random variable, i.e., a measurable function of the state, and let (π, π_0) be a stationary and ergodic continuous-state Markov process (chain). Assume that $E|y| < +\infty$. Then

$$\frac{1}{T} \sum_{t=1}^T y_t \rightarrow E[y] = \int y(s) \pi_0(s) ds$$

with probability 1 with respect to the distribution π_0 .

1.1.9 Example: Asset valuation

Let (\mathbf{P}, π) be a stationary Markov chain with \mathbf{P} that has strictly positive elements, and denote x_t the state at time t . We are interested in the valuation of a dividend process

$\{D_t\}_{t=0}^{\infty}$. We assume that future cash flows are discounted by a constant discount rate $\beta \in (0, 1)$; as we will see later, such a model of discounting would be the outcome of an optimal investment decision made by a risk-neutral agent. In Section 3.1.3, we discuss discounting of risky cash flows in more generality.

The price Q_t of the asset that represents the claim on the dividend process satisfies the recursion

$$Q_t = E_t [\beta (Q_{t+1} + D_{t+1})]. \quad (1.12)$$

The goal is to find a stochastic process $\{Q_t\}_{t=0}^{\infty}$ for the asset price that satisfies this equation. We can roll the equation forward to obtain

$$Q_t = \sum_{j=1}^{T-t} E_t [\beta^j D_{t+j}] + E_t [\beta^{T-t} Q_T],$$

and, under the assumption that

$$\lim_{T \rightarrow \infty} E_t [\beta^{T-t} Q_T] = 0, \quad (1.13)$$

we obtain the **fundamental solution**

$$Q_t = \sum_{j=1}^{\infty} E_t [\beta^j D_{t+j}]. \quad (1.14)$$

To make the model more specific, let us assume that the dividend $D_t = d(x_t) = \bar{d}' x_t$ is a function of the Markov state, and that $\bar{d} \geq 0$ and nonzero. In this case

$$E_t [\beta^j D_{t+j}] = \beta^j \mathbf{P}^j \bar{d}.$$

Denote λ_i the eigenvalues of \mathbf{P} in descending order. Since \mathbf{P} is a transition matrix, $\lambda_1 = 1$, and since \mathbf{P} has strictly positive elements, the Perron–Frobenius theorem implies that $|\lambda_i| < 1$ for $i \geq 2$. Assume for simplicity that the remaining eigenvalues are distinct, so that the associated eigenvectors v_i form a basis of the n -dimensional vector space, and we can write \bar{d} as

$$\bar{d} = c^1 v_1 + c^2 v_2 + \dots + c^n v_n$$

for some projection coefficients c^i . The eigenvector v_1 can be normalized to be strictly positive, and since \bar{d} is positive, it must be that $c^1 > 0$. Indeed, we know from the Perron–Frobenius theorem that v_2 has negative elements. Assume that $c^1 = 0$. Then the contribution of the eigenvector v_2 dominates for large j

$$\mathbf{P}^j \bar{d} \overset{j \rightarrow \infty}{\approx} \lambda_2^j c^2 v_2$$

and hence the conditional expectation of a positive random variable would turn negative, a contradiction.

This implies that

$$E_t \left[\beta^j D_{t+j} \right] = \beta^j \mathbf{P}^j \bar{d} \stackrel{j \rightarrow \infty}{\approx} \beta^j \lambda_1^j c^1 v_1 = \beta^j c^1 v_1,$$

so the value of positive long-horizon payoffs decays at the discount rate β .

We now want to turn to the characterization of solutions $\{Q_t\}_{t=0}^{\infty}$. Assume first that we are looking for a **Markov solution** of the form

$$Q_t = q(x_t) = \bar{q}' x_t.$$

Since the Markov solution is bounded, condition (1.13) is satisfied, and the Markov solution is the fundamental one. In fact, we can solve for it by plugging in the Markov form into the recursive equation (1.12):

$$\bar{q} = \beta \mathbf{P} (\bar{q} + \bar{d}).$$

Since $\beta \mathbf{P}$ has all eigenvalues strictly inside the unit circle, then $I - \beta \mathbf{P}$ is invertible, and we obtain the solution

$$\bar{q} = (I - \beta \mathbf{P})^{-1} \bar{d}.$$

However, can there be other solutions as well? Indeed there can but, as our previous analysis implies, they cannot be functions of the Markov chain. To see this, consider the solution to the following recursive equation

$$Q_t^b = \beta E_t \left[Q_{t+1}^b \right]. \quad (1.15)$$

Conjecture the existence of such a process and define

$$\varepsilon_{t+1} = Q_{t+1}^b - E_t \left[Q_{t+1}^b \right]$$

to be its one-period-ahead expectational error. Then

$$Q_t^b = \beta \left(Q_{t+1}^b - \varepsilon_{t+1} \right),$$

and hence

$$Q_{t+1}^b = \frac{1}{\beta} Q_t^b + \varepsilon_{t+1}.$$

Any stochastic process that satisfies this stochastic difference equation is indeed a solution to (1.15). Observe that from the perspective of valuation, equation (1.15) is a recursive equation for the value of an asset that pays zero cash flows. A positive solution for the price process is called a **bubble**.

Observe that this bubble grows on average at the rate β^{-1} , and violates condition (1.13):

$$\lim_{T \rightarrow \infty} E_t \left[\beta^{T-t} Q_T^b \right] = Q_t^b \neq 0.$$

Let us now combine the fundamental solution and the bubble to define the price process

$$Q_t = \bar{q}'x_t + Q_t^b.$$

This price process $\{Q_t\}_{t=0}^{\infty}$ satisfies the recursive valuation equation (1.12). It is, however, not Markov. The fundamental component $\bar{q}'x_t$ is stationary but the bubble component Q_t^b is not, and hence cannot be written as a function of the Markov state.

1.1.10 Example: Wages in a lake model

In this example, we study the distribution of workers' wages in a version of the McCall (1970) search model. We develop optimal decision-making in this model in Section 2.1. In this example, all transitions between employment states are exogenous.

A worker in the economy can be either employed or unemployed. At the beginning of each period, a previously unemployed worker receives a wage offer with probability λ and accepts it. Also at the beginning of each period, a previously employed worker separates from a job with probability δ .

Employment status distribution

Transitions in the employment status can be represented using a 2-state Markov chain with states e_1, e_2 corresponding to unemployment and employment state $s \in \{u, e\}$, respectively. The transition matrix is given by

$$\mathbf{P} = \begin{bmatrix} 1 - \lambda & \lambda \\ \delta & 1 - \delta \end{bmatrix}.$$

Stationary distributions of the Markov chain are left eigenvectors of \mathbf{P} associated with the unit eigenvalue, normalized to a unit sum, i.e., normalized solutions to $\pi' \mathbf{P} = \pi$. If at least one of the values λ and δ is larger than zero, the unique solution is

$$\pi' = \left(\frac{\delta}{\lambda + \delta}, \frac{\lambda}{\lambda + \delta} \right).$$

Then the stationary Markov chain (\mathbf{P}, π) is ergodic. This means that the stationary cross-sectional distribution of workers' employment statuses can be replicated by drawing a random worker with initial status arbitrarily drawn from π , and following the distribution of her employment status over time. This is true even in the case when either λ or δ are equal to zero.

When $\lambda = \delta = 0$, the model replicates Example 1.4. Any $\pi = (p, 1 - p)'$ for $p \in [0, 1]$ constitutes a stationary distribution. Each individual history has a constant employment status. This means that when $p \in (0, 1)$, the stationary Markov chain (\mathbf{P}, π) is not ergodic, but it is ergodic when $p \in \{0, 1\}$.

Wage distribution

Continuing with the above example, assume that $\lambda, \delta > 0$. Further assume that the wage offer for an unemployed worker is drawn from a distribution with cdf $F(w)$ that has a density $f(w)$ with a full support on $[0, 1]$. Assume that every job offer is accepted by the unemployed worker, even when this is not optimal. When the worker accepts, she stays with the same wage w until she separates into unemployment. An unemployed worker earns $c \geq 0$.

While in this simple model, transitions depend only on employment status, we want to establish the distribution of income. The state in this model is the pair $x = (y, s)$ where $s \in \{u, e\}$ is the employment status, $y = c$ when $s = u$, and $y = w \in [0, 1]$ when $s = e$.

For the unemployed worker, the transition probability is completely characterized by

$$\begin{aligned}\Pi(s' = u \mid s = u) &= 1 - \lambda \\ \pi(y' = w; s' = e \mid s = u) &= \lambda f(w).\end{aligned}$$

The conditional distribution of wage offers conditional on $s' = e$ and $s = u$ therefore has density $f(w)$. On the other hand, when $s = e$, then

$$\begin{aligned}\Pi(s' = u \mid y = w; s = e) &= \delta \\ \Pi(y' = w; s' = e \mid y = w; s = e) &= 1 - \delta.\end{aligned}$$

Using the stationary distribution of the employment statuses, we can infer that the unique stationary distribution Π_0 of (y, s) is

$$(y, s) : \begin{cases} (c, u) & \text{with probability } \Pi_0(c, u) = \frac{\delta}{\lambda + \delta} \\ (w, e) & \text{with density } \pi_0(w, e) = \frac{\lambda}{\lambda + \delta} f(w) \end{cases}$$

The unique stationary distribution of income y is analogous

$$y : \begin{cases} y = c \text{ (unemployed)} & \text{with probability } \Pi_0(0) = \frac{\delta}{\lambda + \delta} \\ y = w \geq 0 \text{ (employed)} & \text{with density } \pi_0(w) = \frac{\lambda}{\lambda + \delta} f(w) \end{cases}$$

Since $\lambda, \delta > 0$, the Markov chain for (y, s) is ergodic. Any set of wages B that has a positive probability under Π_0 and any employment status can be reached with a positive probability from any current state (y, s) , and hence any invariant function has to be constant on the whole support of the stationary distribution.

When $\delta = 0$, there are infinitely many stationary distributions of (y, s) , all characterized by $\Pi_0(c, u) = 0$, but with arbitrary distributions of wages $\pi_0(w)$. Once a worker receives an offer, she stays at the job with wage w forever, so any $\pi_0(w)$ that integrates to one can be supported as the stationary distribution. Under such a distribution, any function $y(w)$ is invariant, since w stays constant and hence $y(w)$ also stays constant. Hence, when $\pi_0(w)$ is nondegenerate, the Markov chain is not ergodic. Only when $\pi_0(w)$ is degenerate and consists of a single mass point, the chain is ergodic.

1.2 Investors' consumption-saving decision problem

We now characterize the decision problem of an investor who can smooth consumption by investing in a variety of assets with potentially stochastic returns. We want to carefully derive the first-order conditions from the sequence problem, paying attention to the appropriate measurability of the individual objects. In order to do so, it is instructive to be explicit about encoding of states. To avoid technical difficulties, we follow the setup from Section A.1 and assume that at every date, one of a finite number S of different states can be realized. Recall that we denoted $s^t = (s_0, s_1, \dots, s_t)$ the history of realized states up to time t , $P(s^t)$ its unconditional probability, and $P(s^{t+j}|s^t)$ the probability of s^{t+j} conditional on a partial history s^t .

There are N available assets, with time- t payoffs $G_t^n(s^t)$ and ex-dividend prices $Q_t^n(s^t)$, $n = 1, \dots, N$. An investor can buy the asset at time t for $Q_t^n(s^t)$, collect payoff $G_{t+1}^n(s^{t+1})$ at time $t + 1$, and subsequently resell the asset for $Q_{t+1}^n(s^{t+1})$, hence earning the return

$$R_{t+1}^n(s^{t+1}) = \frac{Q_{t+1}^n(s^{t+1}) + G_{t+1}^n(s^{t+1})}{Q_t^n(s^t)}. \quad (1.16)$$

We denote $a_t^n(s^t)$ the quantity of asset n bought at time t with history s^t . Buying and selling the assets is unconstrained.

The investor is endowed with expected utility preferences (von Neumann and Morgenstern (1947)) with period utility function $u(c)$ and time-preference parameter $\beta \in [0, 1)$. Every period, the investor also receives an exogenous endowment $y_t(s^t)$.

In order to keep the problem disciplined, we impose, for the purposes of this section, that $Q_t^n(s^t)$, $G_t^n(s^t)$, and $y_t(s^t)$ are bounded stochastic processes. This allows a rigorous characterization of the solution without the need to take into account specific pathological cases with unbounded cash flows or endowments.

With this notation, we can write the consumption-savings problem as

$$\max_{\{c_t(s^t), a_t^n(s^t)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \sum_{s^t} P(s^t|s_0) \beta^t u(c_t(s^t)) \quad (1.17)$$

subject to the sequence of budget constraints

$$c_t(s^t) + \sum_{n=1}^N a_t^n(s^t) Q_t^n(s^t) = \sum_{n=1}^N a_{t-1}^n(s^{t-1}) (Q_t^n(s^t) + G_t^n(s^t)) + y_t(s^t), \quad (1.18)$$

for $t = 1, 2, \dots$, with the initial state s_0 and initial asset positions $a_{-1}^n(s^{-1}) = a_{-1}^n$ given. This set of budget constraints needs to be extended by including a constraint that prevents the agent from running into ever-increasing debt (negative financial wealth), a so-called no-Ponzi condition:

$$\lim_{t \rightarrow \infty} \sum_{n=1}^N a_t^n(s^t) Q_t^n(s^t) \geq -b, \quad (1.19)$$

where \underline{b} is an arbitrarily large number called the debt limit. Along the optimal path, this debt limit will never be reached, but without this constraint, the problem would not be well posed. A more detailed discussion on the role of the no-Ponzi conditions is provided in Section 1.2.2.

The right-hand side of the budget constraint (1.18) indicates the available resources at time t , consisting of the cash flows from the asset positions, proceeds from asset sales, and the exogenous endowment. The timing of the variables indicates that assets being sold at time t have been purchased in quantities $a_{t-1}^n(s^{t-1})$ that have been determined in period $t-1$, i.e., conditional on information contained in the history s^{t-1} . Hence, these asset positions are \mathcal{F}_{t-1} -measurable. The left-hand side represents time- t expenditures, consisting of consumption $c_t(s^t)$ and new asset purchases with quantities $a_t^n(s^t)$. Both these choices are \mathcal{F}_t -measurable. Hence the whole constraint (1.18) is \mathcal{F}_t -measurable.

The notation in (1.17)–(1.18) also reveals the notion how [Debreu \(1959\)](#) and [Arrow \(1964\)](#) thought about the problem of state dependence. Consumption goods in different states constitute distinct goods, hence, given the initial state s^0 , there are S^t distinct goods at time t , one for each s^t . The expected utility operator in (1.17) constitutes a utility aggregator over the consumption bundle $\{c_t(s^t) : t \geq 0, s^t \in \mathcal{S}^t, s_0 \text{ given}\}$ with utility additively separable across states and time.

The setup incorporates a variety of assets. For example, a one-period risk-free bond purchased at time t has $G_{t+1}^n(s^{t+1}) = 1$ and $Q_{t+1}^n(s^{t+1}) = 0$, so that the one-period return, called the risk-free rate, is

$$R_{t+1}^n(s^{t+1}) = \frac{1}{Q_t^n(s^t)}.$$

Observe that this return only depends on time- t information, and is therefore \mathcal{F}_t -measurable.

We can now form the Lagrangian. Notice that (1.18) says that we have one budget constraint at time t for every history s^t . Therefore, we will also have a Lagrange multiplier for every such history, which we denote $\mu_t(s^t)$. The Lagrange multiplier for the set of constraints at time t is also \mathcal{F}_t -measurable. Moreover, as we will see momentarily, it is useful to choose instead as the Lagrange multiplier its scaled version $P(s^t|s_0) \beta^t \mu_t(s^t)$.

Then the Lagrangian is given by

$$\begin{aligned} \mathcal{L}(s_0) = & \sum_{t=0}^{\infty} \sum_{s^t} P(s^t|s_0) \beta^t u(c_t(s^t)) \\ & + \sum_{t=0}^{\infty} \sum_{s^t} P(s^t|s_0) \beta^t \mu_t(s^t) \left[\sum_{n=1}^N a_{t-1}^n(s^{t-1}) (Q_t^n(s^t) + G_t^n(s^t)) + y_t(s^t) \right. \\ & \left. - c_t(s^t) - \sum_{n=1}^N a_t^n(s^t) Q_t^n(s^t) \right] \end{aligned}$$

The Lagrangian omits the no-Ponzi condition, already anticipating that it will be slack along the optimal path. The first-order necessary conditions for optimal consumption and

investment choice can be written as:

$$\begin{aligned} [c_t(s^t)] &: P(s^t|s_0) \beta^t u'(c_t(s^t)) = P(s^t|s_0) \beta^t \mu_t(s^t) \\ [a_t^n(s^t)] &: P(s^t|s_0) \beta^t \mu_t(s^t) Q_t^n(s^t) = \\ &= \sum_{s^{t+1}|s^t} P(s^{t+1}|s_0) \beta^{t+1} \mu_{t+1}(s_{t+1}) \left(Q_{t+1}^n(s^{t+1}) + G_{t+1}^n(s^{t+1}) \right). \end{aligned}$$

Observe that the time- t choice of assets $a_t^n(s^t)$ for a particular s^t determines the asset positions along all paths s^{t+1} that are continuation paths of s^t , i.e., that can be reached from s^t . We denote these paths $s^{t+1}|s^t$. There are S such paths, and the sum on the right-hand side of the first-order conditions sums across them. Simplifying the expressions, we obtain

$$\begin{aligned} [c_t(s^t)] &: u'(c_t(s^t)) = \mu_t(s^t) \tag{1.20} \\ [b_{t+1}(s^t)] &: Q_t^n(s^t) = \beta R \sum_{s^{t+1}|s^t} \frac{P(s^{t+1}|s_0)}{P(s^t|s_0)} \frac{\mu_{t+1}(s^{t+1})}{\mu_t(s^t)} \left(Q_{t+1}^n(s^{t+1}) + G_{t+1}^n(s^{t+1}) \right). \end{aligned}$$

These first-order conditions, together with the sequence of budget constraints, are necessary for an optimum but not sufficient. They assure that the sequence of consumption choices $c_t(s^t)$ is optimal in a relative sense, namely that if $c_t(s^t)$ has been chosen optimally, $c_{t+1}(s^{t+1})$ will also be chosen optimally and so on. In order to also pin down the overall level of the consumption path an additional condition, called the transversality condition, needs to be added that assures that asymptotically, the investor does not oversave resources, in the sense that the value of accumulated financial assets does not explode:

$$\lim_{t \rightarrow \infty} \beta^t \sum_{s^t} P(s^t|s_0) u'(c_t(s^t)) \sum_{n=1}^N a_t^n(s^t) Q_t^n(s^t) = 0.$$

More discussion on the necessity and sufficiency of transversality conditions is provided in Section 1.2.2, in [Kamihigashi \(2003, 2005\)](#), as well as in [Stokey et al. \(1989\)](#), Chapter 4.

Since the ratio of probabilities in (1.20) is the conditional probability

$$\frac{P(s^{t+1}|s_0)}{P(s^t|s_0)} = P(s^{t+1}|s^t),$$

we can rewrite the pair of optimality conditions (1.20) using the conditional expectations operator, and dropping the notation for the histories s^t as

$$\begin{aligned} u'(c_t) &= \mu_t \\ Q_t^n &= E \left[\beta \frac{\mu_{t+1}}{\mu_t} (Q_{t+1}^n + G_{t+1}^n) \mid \mathcal{F}_t \right] = E_t \left[\beta \frac{\mu_{t+1}}{\mu_t} (Q_{t+1}^n + G_{t+1}^n) \right]. \end{aligned}$$

Substituting out the Lagrange multipliers, we obtain the **Euler equation**

$$Q_t^n = E_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} (Q_{t+1}^n + G_{t+1}^n) \right]. \tag{1.21}$$

Similarly, the transversality condition can be written as

$$\lim_{t \rightarrow \infty} \beta^t E_0 \left[u'(c_t) \sum_{n=1}^N a_t^n Q_t^n \right] = 0. \quad (1.22)$$

The Euler equations (1.21) can also be rewritten in returns form using the definition of returns in (1.16) as

$$1 = E_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} R_{t+1}^n \right], \quad n = 1, \dots, N. \quad (1.23)$$

The Euler equations thus constitute restrictions on the joint dynamics of investors' consumption on asset returns in the form of moment conditions. We can therefore use them to estimate preference parameters and test their validity using the generalized method of moments, an approach we take up in Section 1.4.

The Euler equations indicate that from the perspective of an optimizing investor, the price of the asset must be equal to the expected next-period payoff and resale value, adjusted by the marginal rate of substitution. The marginal rate of substitution in this model is commonly called the stochastic discount factor process S_t , with one-period increment

$$\frac{S_{t+1}}{S_t} = \beta \frac{u'(c_{t+1})}{u'(c_t)}.$$

The Euler equations then read

$$1 = E_t \left[\frac{S_{t+1}}{S_t} R_{t+1}^n \right], \quad n = 1, \dots, N. \quad (1.24)$$

This form of asset pricing restrictions emerges in many dynamic environments, including ones in which asset prices are not determined by individual investors maximizing utility from consumption. As we discuss in Section 1.6, absence of arbitrage in finance markets is equivalent to the existence of a strictly positive stochastic discount factor that satisfies restrictions (1.24).

1.2.1 Subjective beliefs

In the consumption-saving problem discussed in the previous section, the investor was endowed with separable preferences over consumption and correct beliefs about the distribution of future states. Such preferences satisfy the [von Neumann and Morgenstern \(1947\)](#) expected utility axioms.

In principle, the investor could be endowed with subjective expected utility axiomatized by [Savage \(1954\)](#), with a probability measure $P^I(s^t)$ replacing $P(s^t)$ throughout the derivation, concluding with Euler equations

$$1 = E_t^I \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} R_{t+1}^n \right], \quad n = 1, \dots, N. \quad (1.25)$$

where $E_t^I[\cdot]$ is the expectations operator under the subjective probability measure P^I . This implies that the restrictions on asset returns are not only determined by investors' marginal rate of substitution but also by their beliefs. This will have profound implications for empirical testability of these restrictions using GMM in Section 1.4.

1.2.2 Transversality conditions vs No-Ponzi conditions

The characterization of the optimal trajectory in the consumption-saving problem involved two types of asymptotic conditions: transversality conditions, and no-Ponzi conditions. While they both involve statements about the asymptotic behavior of dynamic paths, they are completely distinct.

Transversality conditions (TVC) of the type (1.22) are **optimality conditions**, i.e., conditions that the optimal path must satisfy. In finite-horizon problems, they describe how the optimal path must behave as we 'transverse' beyond the planning horizon. In infinite-horizon problems, they describe the asymptotic behavior of the optimal path. The model analyzed in this section involves a free terminal boundary that does not specifically constrain the asymptotic behavior of the optimal path but specific modifications of the transversality conditions can accommodate such additional constraints as well.

In the context of economic models, they assure that the agent does not overaccumulate a resource asymptotically in a way that would leave a strictly positive value of the reward flow on the table in the long run. This is exactly what condition (1.22) expresses—the investor should not overaccumulate financial wealth. This discounted total value, expressed in utils, should converge to zero, otherwise the agent could increase the value function by consuming some of the financial wealth at a finite date t . In fact, condition (1.22) can in some cases be replaced with an inequality (\leq) constraint. [Stokey et al. \(1989\)](#), Chapter 4, provide a pedagogical treatment.

No-Ponzi conditions (NPC) of the type (1.19), on the other hand, are institutional restrictions on the set of feasible paths for the optimization problem to be well-defined. They rule out paths that would imply improvement of agent's objective but would lead to pathological, economically uninteresting solutions, like ever accelerating debt accumulation in an infinite horizon problem. Condition (1.19) explicitly provides a fixed lower bound on negative value of financial assets. Since the stock of financial assets constitutes a valuable resource, the agent should not be allowed to deplete its stock to $-\infty$, even asymptotically. Otherwise, it would be optimal to do so, a result that is not economically interesting. The no-Ponzi condition is thus an **institutional restriction** on the choice set of the agent. Notice that in the case of a variable that represents a negative resource, like the stock of debt in the example below, the appropriate no-Ponzi condition would restrict the variable from above.

In many applications, a sufficient no-Ponzi condition can be represented by bounding the discounted value of the valuable resource from below by zero, while the transversality condition bounds the discounted value of the resources from above by zero. Hence an equality restriction that equates this discounted value with zero would seem to subsume both the TVC and the NPC. But this would be a misinterpretation of their distinct roles.

First, the NPC is imposed on the problem as one of the exogenous institutional constraints, and then, such an appropriately restricted problem implies necessary and sufficient optimality conditions that include the TVC.

We can illustrate these conditions and their roles with a simple deterministic version of the consumption-saving model. Consider an investor who has access to only a single asset, a one-period risk-free bond with a constant interest rate R , or price $Q = R^{-1}$. We denote b_t the amount of borrowing (i.e., a negative investment in the bond) that must be repayed at the beginning of period t . The investor maximizes present discounted utility from consumption

$$\max_{\{b_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c_t)$$

subject to the set of constraints:

$$\begin{aligned} c_t + b_t &= R^{-1}b_{t+1} + y, & \forall t \geq 0 \\ c_t &\geq 0, & \forall t \geq 0 \end{aligned} \quad (1.26)$$

$$\lim_{t \rightarrow \infty} \frac{b_t}{R^t} \leq 0. \quad (1.27)$$

Restriction (1.26) is the budget constraint. The left-hand side denotes time- t expenditures and the right-hand side captures sources of financing, consisting of a fixed endowment y and new borrowing that needs to be repayed at time $t + 1$. Consumption is restricted to be nonnegative.

Restriction (1.27) is the no-Ponzi condition. First observe that debt is an adverse state for the household (more debt is worse), and hence the no-Ponzi condition restricts debt accumulation from above. Second, notice that the bound is not fixed, it merely prevents debt from growing too quickly (at a faster rate than the rate of interest). This is looser than condition (1.19) but still sufficient to exclude pathological paths.

The Euler equations and the transversality condition for this problem are given by

$$\begin{aligned} u'(c_t) &= \beta R u'(c_{t+1}) \\ \lim_{t \rightarrow \infty} -\beta^t u'(c_t) b_t &= 0. \end{aligned}$$

For an exact derivation of the transversality condition, see [Stokey et al. \(1989\)](#), Chapter 4. The transversality condition can then be rewritten using the Euler equations as

$$0 = \lim_{t \rightarrow \infty} -\beta^t u'(c_t) b_t = \lim_{t \rightarrow \infty} -\beta^t \frac{u'(c_0)}{(\beta R)^t} b_t = -u'(c_0) \lim_{t \rightarrow \infty} \frac{b_t}{R^t}.$$

Hence the transversality condition has a structure analogous to the no-Ponzi condition, and it would seem that the no-Ponzi condition is not needed, since the optimal path satisfying the transversality condition already satisfies the no-Ponzi condition (1.27). This intuition is false.

To explain why, consider a candidate optimal path $\{b_{t+1}^*\}_{t=0}^{\infty}$ and the associate con-

sumption path $\{c_t^*\}_{t=0}^\infty$ that satisfy the Euler equation and transversality condition. Now consider a perturbation of the optimal path. Pick a period s , and define $\tilde{c}_s = c_s^* + 1$. For every other period $t \neq s$, keep $\tilde{c}_t = c_t^*$. Notice that this corresponds to borrowing one extra unit of debt in period t , consuming it and rolling it over forever.

Without loss of generality, pick $s = 0$. Notice that under the optimal and modified policy, we have

$$\begin{aligned} c_0^* + b_0^* &= R^{-1}b_1^* + y \\ \tilde{c}_0 + b_0^* &= R^{-1}\tilde{b}_1 + y \end{aligned}$$

subtracting and reorganizing, we get

$$\tilde{b}_1 = b_1^* + R.$$

But rolling over the debt further, we have

$$\tilde{b}_t = b_t^* + R^t. \tag{1.28}$$

Dividing this expression by R^t and taking the limit, we must have

$$\lim_{t \rightarrow \infty} \frac{\tilde{b}_t}{R^t} = \lim_{t \rightarrow \infty} \frac{b_t^*}{R^t} + 1 = 1.$$

Hence the alternative path violates the no-Ponzi scheme condition but at the same time also the transversality condition. But clearly, the policy $\{\tilde{c}_t\}_{t=0}^\infty$ delivers a higher value than $\{c_t^*\}_{t=0}^\infty$, so the latter cannot be optimal. What argument invalidates the use of the transversality condition?

We have argued above that in order for the Euler equations and transversality condition to be sufficient for optimum, some lower bound must be imposed on variables that represent a valuable resource. Since debt b_t is the negative of a valuable resource, it corresponds to an upper bound on b_t , which, in this particular problem, is given by the no-Ponzi condition (1.27). If we restrict feasible paths to those that also satisfy the no-Ponzi condition, then the transversality condition is sufficient and indeed recovers the optimal path.

1.3 Equilibrium in endowment economies

In Section 1.2, we studied the individual decision problem of an investor who takes asset prices as given. We now describe how equilibrium asset prices are determined. We utilize the endowment economy framework analyzed by Lucas (1978).

Time is discrete and infinite, indexed by $t = 0, 1, 2, \dots$. In every period t , one of finitely many states $s_t \in \mathcal{S}$ is realized, with a history of states up to time t denoted $s^t = (s_0, s_1, \dots, s_t)$. The evolution of the states is determined by the data-generating probability measure P , with probabilities over histories $P(s^t)$. The conditional probability of history s^t conditional on partial history s^{t-j} is denoted $P(s^t | s^{t-j})$. Aggregate endowment

in the economy is given by an exogenous process $Y_t = Y(s^t)$.

The representative investor receives individual endowment $y_t = y(s^t)$ and trades a set of trades a set of N assets with prices $Q_t^n = Q^n(x^t)$ and promised cash flows $G_t^n = G^n(x^t)$, $n = 1, \dots, N$. The preferences of the investor are described by a separable utility function $u(c)$ with usual properties, a time-preference parameter $\beta \in [0, 1)$ and a subjective probability measure P^I over the states of the economy. Such subjective expected utility satisfies, for example, the [Savage \(1954\)](#) axioms. The representative investor trades in a competitive market, taking asset prices as given. The decision problem is characterized by the objective function given in (1.17), with optimal behavior analyzed in Section 1.2.

The assets are in fixed supply, given by the vector $\theta = (\theta_1, \dots, \theta_N)$.

Definition 1.8. *A competitive equilibrium in this endowment economy consists of the endowment process $Y(s^t)$, aggregate consumption $C(s^t)$, cash flows $G^n(s^t)$, the price processes $Q^n(s^t)$, and individual allocations $c(s^t)$, $a^n(s^t)$, $n = 1, \dots, N$, such that:*

1. *given prices $Q^n(s^t)$, cash flows $G^n(s^t)$, $n = 1, \dots, n$, and individual endowment $y(s^t)$, the investor chooses consumption $c(s^t)$ and portfolio allocation $a^n(s^t)$, $n = 1, \dots, n$, that solve the utility maximization problem (1.17)–(1.19),*
2. *the individual investor is representative, so that individual and aggregate variables*

$$c(s^t) = C(s^t) \quad y(s^t) = Y(s^t),$$

3. *markets clear*

$$\begin{aligned} C(s^t) &= \sum_{n=1}^N \theta_n G^n(s^t) + Y(s^t) \\ b^n(s^t) &= \theta_n \quad n = 1, \dots, N. \end{aligned}$$

The equilibrium is a rational expectations equilibrium if investor's subjective probability measure P^I coincides with the data-generating probability measure P .

In the equilibrium in the endowment economy, cash flows G_t^n and endowment Y_t are exogenous, while asset prices Q_t^n are endogenously determined.

The representative investor assumption builds on the idea that the given investor is a representative member of a class of investors who share the same preferences and environment. As a consequence, individual and aggregate variables (consumption and endowment) coincide. This is analogous to the “little k , big K ” argument in [Ljungqvist and Sargent \(2018\)](#), Chapters 7 and 12.

This also implies that we can also replace the stochastic discount factor in Euler equations (1.23), or in their subjective belief version (1.25), respectively, with

$$1 = E_t^I \left[\beta \frac{u'(C_{t+1})}{u'(C_t)} R_{t+1}^n \right], \quad n = 1, \dots, N \quad (1.29)$$

where the stochastic discount factor now depends on the aggregate consumption process:

$$s_{t+1} = \frac{S_{t+1}}{S_t} = \beta \frac{u'(C_{t+1})}{u'(C_t)}.$$

The definition of the equilibrium also clarifies the distinct roles the data-generating measure P and investor's beliefs P^I play in the determination of the equilibrium. The data-generating measure P describes the evolution of the state s_t over time. On the other hand, given a particular state with history s^t , the equilibrium asset prices $Q^n(s^t)$ are determined by investor's optimality conditions (1.29) that involves the subjective belief P^I . The mapping from a particular history s^t to asset prices $Q^n(s^t)$ therefore depends only on the subjective belief P^I , not on the data-generating measure. The return on the asset

$$R_{t+1}^n = \frac{Q^n(s^{t+1}) + G^n(s^{t+1})}{Q^n(s^t)}$$

then depends on both measures in the following sense. The values $Q^n(s^t)$ and $Q^n(s^{t+1})$ as functions of the histories are determined by the subjective belief P^I . On the other hand, the distribution of s^{t+1} conditional on s^t is determined by the data-generating measure, $P(s^{t+1}|s^t)$.

Under the rational expectations assumption (Muth (1961), Lucas (1972)), the subjective belief P^I and the data-generating probability measure P coincide. From now on, we impose the rational expectations assumptions unless explicitly noted otherwise.

1.3.1 Equilibrium asset prices

The endowment economy setup allows for assets with flexibly specified cash flows. For example, in the case of a stock, the payoff G_t is the dividend stream, and Q_t the ex-dividend price. The asset return then is

$$R_{t+1} = \frac{Q_{t+1} + G_{t+1}}{Q_t}$$

with conditional expected return $E_t[R_{t+1}]$.

In the case of a risk-free bond with maturity T with time- t price $Q_t^{[T]}$, the cash flow stream G_t consists of a sequence of coupons in periods $t+1, \dots, t+T$, and a principal which we can normalize to 1. We can therefore set the price right before the principal is paid off to $Q_{t+T}^{[0]} = 1$, and then the bond prices satisfy the recursion

$$Q_{t+j}^{[T-j]} = E_{t+j} \left[\frac{S_{t+j+1}}{S_{t+j}} \left(Q_{t+j+1}^{[T-j-1]} + G_{t+j+1} \right) \right].$$

In the case of zero-coupon bonds, the coupons $G_t = 0$ and the payoffs consist only of the principal at maturity. The yield to maturity is the average return the bond earns between today and its maturity

$$y_t^{[T]} = -\frac{1}{T} \log Q_t^{[T]}.$$

In the case of the one-period zero-coupon risk-free bond, the price is

$$Q_t^{[1]} = E_t \left[\frac{S_{t+1}}{S_t} \right]$$

and the return, called the risk-free rate, is the reciprocal of the expected value in the increment of the SDF:

$$R_{t+1}^f = \frac{Q_{t+1}^{[0]}}{Q_t^{[1]}} = \left(Q_t^{[1]} \right)^{-1} = \left(E_t \left[\frac{S_{t+1}}{S_t} \right] \right)^{-1}.$$

Observe that this return is known at time t , i.e., is in the time- t information set.

An excess return is the difference between two returns, often between a risky and a risk-free one

$$R_{t+1}^e = R_{t+1} - R_{t+1}^f.$$

Subtracting the two Euler equations for returns R_{t+1} and R_{t+1}^f implies that excess returns satisfy

$$0 = E_t \left[\frac{S_{t+1}}{S_t} R_{t+1}^e \right]. \quad (1.30)$$

The risk premium is the expected excess return between two assets, $E_t [R_{t+1}^e]$.

The Euler equations (1.21) define asset prices recursively. and so we can iterate the recursive equation forward:

$$\begin{aligned} Q_t &= E_t \left[\frac{S_{t+1}}{S_t} G_{t+1} \right] + E_t \left[\frac{S_{t+1}}{S_t} Q_{t+1} \right] = E_t \left[\frac{S_{t+1}}{S_t} G_{t+1} + \frac{S_{t+2}}{S_t} G_{t+1} \right] + E_t \left[\frac{S_{t+2}}{S_t} Q_{t+2} \right] \\ &= \sum_{j=1}^T E_t \left[\frac{S_{t+j}}{S_t} G_{t+j} \right] + E_t \left[\frac{S_{t+T}}{S_t} Q_{t+T} \right]. \end{aligned}$$

Taking the limit of this iteration as $T \rightarrow \infty$ yields

$$Q_t = \underbrace{\lim_{T \rightarrow \infty} \sum_{j=1}^T E_t \left[\frac{S_{t+j}}{S_t} G_{t+j} \right]}_{\text{fundamental value}} + \underbrace{\lim_{T \rightarrow \infty} E_t \left[\frac{S_{t+T}}{S_t} Q_{t+T} \right]}_{\text{rational bubble}}. \quad (1.31)$$

The first term on the right-hand side, representing the present discounted value of future cash flows, is called the fundamental value of the asset. The second term, the asymptotic discounted value of the asset, is called a rational bubble. When rational bubble exist, the current asset value Q_t can be strictly positive even when the asset does not pay any cash flows, $G_{t+j} = 0, j > 0$.

Optimal behavior of individual investors described in the consumption-saving model in Section 1.2 does not preclude the existence of rational bubbles. Investors are willing to buy an asset whose price satisfies equation (1.31) at a positive price Q_t because they believe they will be able to resell this asset in the future at a positive price Q_{t+T} .

Whether bubbles can emerge in a model will therefore depend on the structure of the model. Rational bubbles can emerge in specific models with market structure in which assets can have specific role in providing liquidity, insuring idiosyncratic risk (Bewley (1980)) or facilitating intergenerational trade (as in the overlapping generations economy in Tirole (1985)). Money is an example of such an asset that is a bubble. However, rational expectations equilibria put strong discipline on when bubble can emerge (Santos and Woodford (1997)), since equilibrium restrictions often imply transversality conditions or no-bubble conditions in the form

$$\lim_{T \rightarrow \infty} E_t \left[\frac{S_{t+T}}{S_t} Q_{t+T} \right] = 0.$$

In the representative agent economy outlined in this section, bubbles indeed cannot emerge in equilibrium.

1.3.2 Risk premia as compensation for risk

Differences in expected returns between assets express differential compensation for risk associated with returns on these assets. Recall the expression for the covariance between two random variables X and Y

$$\text{Cov}(X, Y) = E[X]E[Y] - E[XY]. \quad (1.32)$$

Take two returns, for example, a risky return R_{t+1}^n on asset n and the risk-free rate R_{t+1}^f . Since both these returns must satisfy the pricing Euler equation, the excess return $R_{t+1}^{n,e} = R_{t+1}^n - R_{t+1}^f$ satisfies

$$0 = E_t \left[\frac{S_{t+1}}{S_t} \left(R_{t+1}^n - R_{t+1}^f \right) \right].$$

Using the covariance formula with time- t conditioning, we can rewrite this restriction as

$$0 = E_t \left[\frac{S_{t+1}}{S_t} \right] E_t \left[R_{t+1}^n - R_{t+1}^f \right] + \text{Cov}_t \left(\frac{S_{t+1}}{S_t}, R_{t+1}^n - R_{t+1}^f \right).$$

Reorganizing this expression, and noticing that

$$\left(E_t \left[\frac{S_{t+1}}{S_t} \right] \right)^{-1} = R_{t+1}^f$$

we obtain the relationship between expected excess returns and covariances of these returns with the stochastic discount factor

$$E_t \left[R_{t+1}^n - R_{t+1}^f \right] = -R_{t+1}^f \text{Cov}_t \left(\frac{S_{t+1}}{S_t}, R_{t+1}^n - R_{t+1}^f \right). \quad (1.33)$$

These expected excess returns are called risk premia, since they represent required compensation for holding risky assets. In order to interpret the covariance, assume the case of

the CRRA preference SDF

$$\frac{S_{t+1}}{S_t} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma}.$$

In bad state realizations at time $t + 1$, consumption C_{t+1} is low, which implies a high marginal rate of substitution, i.e., a high realization of the SDF. The situation is reversed in good states of the economy. Then an asset that delivers low access returns $R_{t+1}^n - R_{t+1}^f$ in bad states of the economy and high returns in good states will have a negative covariance with the SDF. Such an asset is risky, since it delivers high returns in states of the world in which consumption is not valuable, and vice versa. Investors are then compensated for holding such a risky asset with a higher expected return. Observe that what matters is the covariance with the SDF—an asset with volatile returns should earn a zero risk premium if these returns are uncorrelated with the SDF.

Equation (1.33) lends itself to cross-sectional and time-series tests. In the cross-section, assets with returns that covary more negatively with the stochastic discount factor should earn higher risk premia. In the time series, periods of time in which the conditional covariance is more negative should be associated with higher conditional risk premia.

1.4 Testing Euler equations using generalized method moments

In order to fully specify and solve the endowment economy in Section 1.3, we need to parameterize the data-generating probability measure, the mapping between the states and aggregate endowment, the full set of traded assets and their cash flows, and the preferences of the representative agent. The model then predicts the dynamics of asset prices that can be compared to the data using formal empirical tests. From a practical standpoint, the structure of these assumptions would need to be sufficiently tractable to keep the computation of the model feasible.

Specifying all these features can be a daunting task. In order to test the model or to identify a specific subset of parameters that are of our interest, we can perhaps resort to a more modest goal, and, in the words of Lars Peter Hansen, “do something without doing everything” (Hansen (2014)). Optimality conditions in the form of Euler equations

$$0 = E_t \left[\frac{S_{t+1}}{S_t} R_{t+1}^n - 1 \right] \quad (1.34)$$

are moment conditions that can be tested using the generalized method of moments (GMM), developed by Hansen (1982), initially implemented in the asset pricing context by Hansen and Singleton (1982). If we specify a model of the stochastic discount factor that can be empirically implemented using observable quantities, then the dynamics of data on the stochastic discount factor and asset returns are jointly restricted by the Euler equations, which can be used to estimate parameters of the stochastic discount factor and to expose the model to a formal test. The asset pricing literature has widely used GMM to estimate models of preferences as well as rejecting large classes of preferences as empirically implausible.

A second reason why GMM estimation may be advantageous is the ability to avoid misspecifications emerging from parts of the model that we may have less confidence in. If the full model is correctly specified, then utilizing only the moment conditions may neglect information relative to using the full likelihood function of the model, and GMM may be less efficient. However, there may be good reasons to trust specific Euler equations more relative to some other details of the model incorporated in the likelihood, and so more substantial concerns that estimation based on likelihood maximization is misspecified.

Since the Euler equations constitute only a subset of the predictions from the original model, the estimation and testing procedure will not be able to determine all parameters of the model. This approach is called partial identification of the model.

As mentioned above, the idea underlying the moment condition test is to take unconditional expectations of equation (1.34) to obtain

$$0 = E \left[\frac{S_{t+1}}{S_t} R_{t+1}^n - 1 \right] \quad (1.35)$$

and then use time series data for the SDF and returns to replace the theoretical moment with its empirical counterpart

$$0 = \frac{1}{T} \sum_{t=0}^{T-1} \frac{S_{t+1}}{S_t} R_{t+1}^n - 1.$$

A law of large numbers implies that the empirical moment converges to the theoretical one as $T \rightarrow \infty$ if the model is correctly specified.

The unconditional moment restrictions (1.35) are weaker than their conditional counterparts (1.34), they only test whether the Euler equation holds on average, and will not be able to uncover conditional violations. However, what if we are concerned that the conditional equation does not hold? In that case, we should be able to predict *when* it does not hold. We denote such predictor variables, or instruments, z_t^k , $k = 1, \dots, K$. Instruments are stochastic processes that are suspected to vary systematically with potential violations of the conditional Euler equation. For example, z_t can be a cyclical variable that can be used to test for systematic violations of the conditional Euler equation over the business cycle. In this way, instruments incorporate conditioning information in the moment conditions.

Multiplying the conditional Euler equation (1.34) with the instrument z_t^k and taking unconditional expectations, we obtain

$$0 = E \left[z_t^k E_t \left[\frac{S_{t+1}}{S_t} R_{t+1}^n - 1 \right] \right] = E \left[z_t^k \left(\frac{S_{t+1}}{S_t} R_{t+1}^n - 1 \right) \right].$$

Again, this is an unconditional moment that can be replaced with a time series average. With $z_t^k = 1$, we obtain the original unconditional moment as a special case.

From the perspective of the cross-sectional and time-series interpretation of equation (1.33), including unconditional equations of the form (1.35) for additional assets n increases the cross-sectional dimension of the restrictions, while including instruments z_t^k for a given

return strengthens the information content of the time-series dimension.

To make the problem more concrete, let us assume that the representative investor is endowed with constant relative risk aversion (CRRA) preferences

$$\begin{aligned} u(c) &= \frac{c^{1-\gamma} - 1}{1-\gamma}, & 1 \neq \gamma > 0 \\ u(c) &= \log(c), & \gamma = 1. \end{aligned} \quad (1.36)$$

Then the SDF is given by

$$\frac{S_{t+1}}{S_t} = \beta \frac{u'(c_{t+1})}{u'(c_t)} = \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\gamma} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma}$$

where the last equality that equates individual and aggregate consumption follows from the representative agent assumption. We then have Euler equations

$$0 = E_t \left[\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{t+1}^n - 1 \right]$$

for a range of asset returns $n = 1, \dots, N$. Further moments can be obtained by including instruments z_t^k . We then have available data

$$X_{t+1} = \left(C_{t+1}/C_t, R_{t+1}^1, \dots, R_{t+1}^N, z_t^1, \dots, z_t^K \right)$$

for $t = 0, \dots, T-1$, and want to estimate the vector of unknown parameters $\theta = (\beta, \gamma)$ which provides the best fit to the Euler equations, and subsequently test whether the restrictions in the form of Euler equations at the estimated parameter values are consistent with the data. We have in total M moment conditions for various combinations of returns and instruments that can be written by defining moment functions

$$f_m(X_{t+1}; \theta) = z_t^k \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{t+1}^n - 1 \right), \quad m = 1, \dots, M.$$

These moment functions can be stacked into a vector

$$f(X_{t+1}; \theta) = (f_1(X_{t+1}; \theta), \dots, f_M(X_{t+1}; \theta))'$$

and the theoretical moment conditions can be written as an $M \times 1$ vector condition

$$0 = E[f(X_{t+1}; \theta)]. \quad (1.37)$$

These Euler equation restrictions make asset pricing applications a particularly fruitful area for the application of the GMM estimation technique. A seminal implementation of this procedure is provided in [Hansen and Singleton \(1982\)](#).

1.4.1 Role of rational expectations in testing Euler equations

As mentioned above, the idea underlying the moment condition test relies on replacing unconditional moments with their time-series counterparts. This was possible because in the Euler equation (1.35), the rational expectations assumption imposed that the investor has correct beliefs that coincide with the data-generating measure.

However, if the investor is endowed with subjective beliefs P^I , the relevant Euler equation implied by the consumption-saving problem is

$$1 = E_t^I \left[\frac{S_{t+1}}{S_t} R_{t+1}^n \right],$$

involving expectations under P^I . The time-series average using data drawn from the data-generating measure will not approximate this moment. More specifically, the conditional Euler equation is

$$\begin{aligned} 1 &= \sum_{s^{t+1}|s^t} P^I(s^{t+1}|s^t) \frac{S(s^{t+1})}{S(s^t)} R^n(s^{t+1}) \\ &= \sum_{s^{t+1}|s^t} P(s^{t+1}|s^t) \frac{P^I(s^{t+1}|s^t)}{P(s^{t+1}|s^t)} \frac{S(s^{t+1})}{S(s^t)} R^n(s^{t+1}) \end{aligned}$$

Denoting

$$m_{t+1} = m(s^{t+1}|s^t) = \frac{P^I(s^{t+1}|s^t)}{P(s^{t+1}|s^t)}$$

we have the Euler equation with involving an expectation under the data-generating measure

$$1 = E_t \left[m_{t+1} \frac{S_{t+1}}{S_t} R_{t+1}^n \right].$$

The quantity m_{t+1} is called the Radon–Nikodým derivative, or the belief ratio, and represents deviations of the subjective belief from the data-generating measure. Hence, in order to expose the Euler equation to a GMM test in the subjective belief environment, we also need to find an empirical counterpart or model for the belief ratio m_{t+1} .

1.4.2 Generalized method of moments

The generalized method of moments characterizes the statistical uncertainty associated with the estimation of an r -dimensional parameter vector θ using an M -dimensional vector of moment conditions

$$0 = E[f(X_{t+1}; \theta)]. \quad (1.38)$$

Specifically, we want to ask how to find the parameter vector θ that yields the best fit to the moment conditions, and how to conduct inference using sample data about the validity of these theoretical restrictions. For the latter, question, we need to establish a distribution theory of how errors in the moment conditions should behave under the null hypothesis

that the moment conditions are correct.

The method was formally derived in Hansen (1982), and an accessible summary with a special attention to time-series applications is provided in Hansen (2001, 2008). The whole Number 4 of Volume 20 of the Journal of Business and Economic Statistics is devoted to a symposium on various areas of application of GMM.

Denote θ_0 the true parameter value, in the sense that it is the unique solution to $0 = E[f(X_{t+1}; \theta)]$. This assumption on a unique solution is called the identification assumption. The solution to the vector equation (1.38) can be equivalently found by solving

$$\theta_0 = \min_{\theta} E[f(X_{t+1}; \theta)]' W E[f(X_{t+1}; \theta)] \quad (1.39)$$

where W is an arbitrary positive definite weighting matrix. The positive definiteness of W implies that the right-hand side of (1.39) is strictly positive if and only if $E[f(X_{t+1}; \theta)]$ is nonzero.

In a finite data sample, we replace the theoretical moment $E[f(X_{t+1}; \theta)]$ with its sample average

$$g_T(\theta) = \frac{1}{T} \sum_{t=0}^{T-1} f(X_{t+1}; \theta)$$

and solve

$$\hat{\theta}_T = \min_{\theta} g_T(\theta)' W g_T(\theta). \quad (1.40)$$

A law of large numbers implies that under appropriate conditions, $g_T(\theta_0) \rightarrow f(x_{t+1}; \theta_0)$, and hence $\hat{\theta}_T \rightarrow \theta_0$.

In principle, any positive definite matrix W is theoretically sufficient for identification. However, specific choices of W have appealing properties. Consider the long-run covariance matrix

$$V = \sum_{j=-\infty}^{\infty} E[f(X_{t+1}; \theta_0) f(X_{t+1+j}; \theta_0)'] . \quad (1.41)$$

This covariance matrix takes into account not only contemporaneous comovement between the components of the moment function but also intertemporal dependence of the data across all horizons. The reason why this expression is a covariance is that at the true parameter value, $E[f(X_{t+1}; \theta_0)] = 0$, so that the last term in the covariance formula (1.32) is zero.

Hansen (1982) shows that when we have more moment conditions than estimated parameters, we can construct an overidentification test based on the statistic

$$T g_T(\theta_0)' V^{-1} g_T(\theta_0) \rightarrow \chi^2(M - r) \quad (1.42)$$

where M is the number of moment conditions and r is the number of estimated parameters. Choosing $W = V^{-1}$, which is asymptotically efficient, intuitively downweights the contribution of the components of the moment function that have a lot of variability, and are therefore less informative. Expression (1.42) provides a formal econometric test of the

hypothesis $\hat{\theta}_T = \theta_0$.

Computation of the weighting matrix

There is a range of issues that need to be considered when computing the long-run covariance matrix V .

First, expression (1.41) involves an infinite sum of autocovariances that accounts for temporal dependence of the data. With a finite data sample for periods $t = 0, \dots, T - 1$, we can approximate each term as

$$E \left[f(X_{t+1}; \theta_0) f(X_{t+1+j}; \theta_0)' \right] \approx \frac{1}{T - |j|} \sum_{i=\max(0, -j)}^{T-1-\max(0, j)} f(X_{t+1+i}; \theta_0) f(X_{t+1+j+i}; \theta_0)'.$$

which indicates that j is bounded by available data. Without imposing additional structure on the temporal dependence, we need to truncate the sum in (1.41) to obtain

$$V = \sum_{j=-L}^L E \left[f(X_{t+1}; \theta_0) f(X_{t+1+j}; \theta_0)' \right]$$

which is equivalent to assuming that the data are not dependent if they are more than L periods apart. For iid data draws, we can then choose $L = 0$. An alternative is to impose a particular theoretical structure on the temporal dependence and then derive a theoretical adjustment for the estimator of V , like in the frequently used autocorrelation- and heteroskedasticity-consistent estimator of [Newey and West \(1987\)](#).

Fortunately, the fact that the moment function was generated from conditional moment restrictions helps simplify the autocovariance significantly. Notice that for $j \geq 1$, we have

$$E \left[f(X_{t+1}; \theta_0) f(X_{t+1+j}; \theta_0)' \right] = E \left[f(X_{t+1}; \theta_0) E_{t+j} \left[f(X_{t+1+j}; \theta_0)' \right] \right] = 0$$

where the first equality uses the law of iterated expectations, and the second equality uses the fact that $E_{t+j} \left[f(X_{t+1+j}; \theta_0)' \right] = 0$ is implied by the conditional Euler equation. In other words, while the data X_{t+1} can, and in fact in most cases will be temporally dependent, the Euler equation errors $f(X_{t+1}; \theta_0)$ will not. It is also worth remembering that this conclusion is based on the assumption that all data X_{t+1} are in the investor's time- $t + 1$ information set, in order to be able to apply the law of iterated expectations in the equation above. The covariance matrix then simplifies to

$$V = E \left[f(X_{t+1}; \theta_0) f(X_{t+1}; \theta_0)' \right] \approx \frac{1}{T} \sum_{t=0}^{T-1} f(X_{t+1}; \theta_0) f(X_{t+1}; \theta_0)'.$$

Second, the theoretical V is a function of the true parameter θ_0 , which is a priori unknown. [Hansen et al. \(1996\)](#) propose to deal with this issue using the so-called continu-

ously updated GMM estimator

$$\hat{\theta}_T = \min_{\theta} g_T(\theta)' V_T(\theta)^{-1} g_T(\theta)$$

where

$$V_T(\theta) = \frac{1}{T} \sum_{t=0}^{T-1} f(X_{t+1}; \theta) f(X_{t+1}; \theta)' \quad (1.43)$$

In this procedure, we search across the parameter space while simultaneously adjusting the weighting matrix. Hansen et al. (1996) report advantageous properties of this procedure in practical situations.

However, a two-step procedure is also asymptotically valid because V in (1.42) can be replaced by a consistent estimator. So we can proceed as follows. First, compute the first-stage estimate $\hat{\theta}_T$ by minimizing (1.40) for some positive definite weighting matrix W . Then compute an estimate of the covariance matrix $V(\hat{\theta}_T)$ using (1.43). This estimator evaluates the Euler equation errors $f(X_{t+1}; \hat{\theta}_T)$ at the parameter value from the first stage. Finally, we can evaluate the left-hand-side of (1.42) using $\theta_0 \approx \hat{\theta}_T$ and $V \approx V(\hat{\theta}_T)$.

In principle, we could also use the estimated $V(\hat{\theta}_T)$ as a new weighting matrix W in the minimization (1.40) and obtain a new estimate $\hat{\theta}_T$, and perhaps iterate on this procedure. These alternative strategies should be evaluated carefully based on the properties of the application and available data.

Finally, there are additional considerations that need to be taken into account, despite the fact that using $V = V(\theta_0)$ as the weighting matrix is asymptotically efficient in a correctly specified model. First, since θ_0 must be estimated using a finite sample of data, we can at best use $V(\hat{\theta}_T)$. Second, the Euler equation errors across moment conditions can be strongly cross-correlated in a way that makes the calculation of $V(\hat{\theta}_T)^{-1}$ fragile, putting excessive emphasis on a small number of moments. This issue is further magnified when potential misspecifications are present. From an empirical standpoint, asset pricing applications are often prone to such fragility, and the literature employed a range of approaches how to modify the weighting matrix to deal with these issues, trading off theoretical efficiency for robustness.

1.5 Valuation in finite-state Markov chain economies

We now provide a tractable implementation of the Lucas (1978) endowment economy introduced in Section 1.3. A specific application of this framework involves one of the seminal papers on the so-called equity premium puzzle by Mehra and Prescott (1985).

Fully solving the model will allow us to completely characterize its predictions, summarized in the likelihood function associated with the model. This is a notably different approach from the GMM-based estimation outlined in Section 1.4. The GMM-based estimation based on the test of Euler equations required very little structure, beyond regularity conditions that guaranteed that the estimator is well-behaved. The application only required data on consumption and returns, and an assumption on the form of the stochastic

discount factor. The Euler equation restrictions then allowed the estimation of the model parameters and an overidentification test of the restrictions. At the same time, much of the structure of the model, such as the parametric specification of the probability distributions, has been put aside.

Here, we take an opposite approach and fully specify the structure of the economy that will endogenize asset prices and yield predictions for the distribution of asset returns. Hansen and Singleton (1983) use an intermediate approach, in which they rely only on Euler equation tests but prespecify exogenous parametric distributions for the joint behavior of consumption and returns.

As in Section 1.3, time is discrete and infinite, $t = 0, 1, 2, \dots$. To impose a tractable structure on the uncertainty, we assume that the probability measure P is characterized by a time-invariant n -state Markov state x_t with transition matrix \mathbf{P} and initial distribution π_0 . For simplicity, we assume that all elements of \mathbf{P} are strictly positive, which implies that the Markov chain has a unique stationary distribution and is ergodic, with details provided in Section 1.3.

With this structure, the abstract state s_t from Section 1.3 with histories s^t is replaced by the Markov state x_t . The Markov structure implies that

$$P(x^{t+1}|x^t) = P(x_{t+1}|x_t)$$

with $P(x_{t+1} = e_j|x_t = e_i) = \mathbf{P}_{ij}$. The probability of a particular history is then given by chaining the transition probabilities

$$P(x^t|x_0) = P(x_t|x_{t-1})P(x_{t-1}|x_{t-2}) \dots P(x_1|x_0)$$

while the probability of being in a particular state x_t conditional on x_0 , irrespective of the particular history that lead to x_t , is

$$P(x_t = e_j|x_0 = e_i) = [\mathbf{P}]_{ij}^t.$$

Observe that $P(x^t|x_0)$ and $P(x_t|x_0)$ are distinct objects, where the latter integrates across all possible intermediate histories (x_1, \dots, x_{t-1}) that lead to x_t .

One possibility how to proceed is to specify aggregate endowment and priced cash flows to be functions of the Markov state. This imposes a stationarity property on the economy because the endowment and cash flows can only take finitely many values and are thus bounded. This is rather restrictive from the perspective of plausible asset pricing models, since empirically, consumption and dividends grow stochastically over time. We therefore allow ourselves more generality, and assume a model with stationary *growth rates* of consumption and cash flows.

In particular, we assume that the growth rate of aggregate endowment C_t is stationary and given by

$$\log C_{t+1} - \log C_t = g_C(x_t, x_{t+1}).$$

This allows for a more general specification than assuming that the consumption *level* is

stationary, $C_t = C(x_t)$, since in this case, we can also write

$$\log C_{t+1} - \log C_t = \log C(x_{t+1}) - \log C(x_t) = g_C(x_t, x_{t+1}).$$

Since x_t is an n -state Markov chain, the function $g_C(x_t, x_{t+1})$ can be expressed using an $n \times n$ matrix. Specifically, we construct an $n \times n$ matrix Γ_C with elements

$$[\Gamma_C]_{ij} = \exp(g_C(x_t = e_i, x_{t+1} = e_j))$$

where we first exponentiated the growth rates for subsequent convenience. For example, in the case of stationary consumption $C_t = C(x_t)$, we have

$$[\Gamma_C]_{ij} = \frac{C(x_{t+1} = e_j)}{C(x_t = e_i)}.$$

We impose the same structure on the stochastic discount factor. In particular, we restrict ourselves to models of the SDF S_t that satisfy

$$\log S_{t+1} - \log S_t = g_S(x_t, x_{t+1}).$$

This stationary specification of the growth rate of the SDF encompasses a very wide class of asset pricing models from the literature. As with consumption growth, we denote Γ_S the $n \times n$ matrix with elements

$$[\Gamma_S]_{ij} = \exp(g_S(x_t = e_i, x_{t+1} = e_j)).$$

Under the assumption that the representative investor is endowed with CRRA preferences (1.36), we have

$$\frac{S_{t+1}}{S_t} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma}$$

and hence

$$\log S_{t+1} - \log S_t = g_S(x_t, x_{t+1}) = \log \beta - \gamma g_C(x_t, x_{t+1}).$$

We also impose that the agent is endowed with correct beliefs, so that investor's subjective probability measure coincides with the data-generating measure.

1.5.1 Asset pricing formulas

Finally, priced cash flows can be stationary, $G_t = G(x_t)$, have stationary growth rates

$$\log G_{t+1} - \log G_t = g_G(x_t, x_{t+1}),$$

or have a different structure like in the case of finite-maturity bonds. We cover the specific implementation of the valuation of these cash flows separately.

The central observation is that asset prices inherit the Markov structure of cash flows and the SDF.

Consider first an asset that pays $G_{t+T} = G(x_{t+T})$ at a single given time $t + T$, and denote $Q_t^{[T]}$ its time- t price, the T standing for time to maturity. An example of such an asset is a zero-coupon bond with $G_{t+T} = 1$. Then for $T = 1$, the price is given by

$$Q_t^{[1]} = Q^{[1]}(x_t) = E_t [\exp(g_S(x_t, x_{t+1})) G(x_{t+1})],$$

where $E_t[\cdot] = E[\cdot|x_t]$. The fact that the price is only a function of the Markov state, $Q_t^{[1]} = Q^{[1]}(x_t)$, follows from the observation that the quantity in the expectation on the right-hand side is only a function of x_t and x_{t+1} , and the expectation integrates over x_{t+1} .

Prices of longer-horizon asset can then be computed iteratively as

$$Q^{[T+1]}(x_t) = E_t [\exp(g_S(x_t, x_{t+1})) Q^{[T]}(x_{t+1})], \quad T = 1, 2, \dots \quad (1.44)$$

Since the prices only depend on the current realization of the n -state Markov chain, they can be encoded as $n \times 1$ vectors. In particular, denote $\mathbf{q}^{[T]}$ the state-dependent price of the maturity- T asset, with elements

$$[\mathbf{q}^{[T]}]_i = Q^{[T]}(x_t = e_i).$$

Using matrix algebra, the recursive equation (1.44) can be expressed as

$$[\mathbf{q}^{[T+1]}]_i = Q^{[T+1]}(x_t = e_i) = \sum_{j=1}^n [\mathbf{P}]_{ij} [\Gamma_S]_{ij} [\mathbf{q}^{[T]}]_j$$

which can be written in compact vector form as

$$\mathbf{q}^{[T+1]} = (\mathbf{P} * \Gamma_S) \mathbf{q}^{[T]}$$

where $\mathbf{P} * \Gamma_S$ is the element-wise multiplication of the two matrices. If the investor was endowed with subjective beliefs represented by a transition matrix \mathbf{P}^l , then \mathbf{P}^l would replace \mathbf{P} in the above formula.

Let us now instead consider the price of an infinite-horizon asset that pays a stationary cash flow $G_t = G(x_t)$ in every period. Then we can value this asset recursively as

$$Q_t = E_t \left[\frac{S_{t+1}}{S_t} (G_{t+1} + Q_{t+1}) \right]. \quad (1.45)$$

In the matrix formulation, the stationary cash flow can be represented by a vector \mathbf{g} with elements

$$\mathbf{g}_i = G(x_t = e_i)$$

and the price Q_t again as a vector \mathbf{q} . We then have the valuation formula

$$\mathbf{q}_i = Q(x_t = e_i) = \sum_{j=1}^n [\mathbf{P}]_{ij} [\Gamma_S]_{ij} (\mathbf{g}_j + \mathbf{q}_j),$$

or, in matrix form

$$\mathbf{q} = (\mathbf{P} * \Gamma_S) (\mathbf{g} + \mathbf{q}).$$

This is a linear system that can be solved for \mathbf{q} to obtain

$$\begin{aligned} [I - (\mathbf{P} * \Gamma_S)] \mathbf{q} &= (\mathbf{P} * \Gamma_S) \mathbf{g} \\ \mathbf{q} &= [I - (\mathbf{P} * \Gamma_S)]^{-1} (\mathbf{P} * \Gamma_S) \mathbf{g} \end{aligned}$$

This formula is a matrix generalization of the geometric sum formula and is valid when the matrix $\mathbf{P} * \Gamma_S$ is stable, in the sense that it has all eigenvalues strictly inside the unit circle. Alternatively, the solution could be obtained as a limit of iterations on the vector $\mathbf{q}^{(n)}$

$$\mathbf{q}^{(n+1)} = (\mathbf{P} * \Gamma_S) (\mathbf{g} + \mathbf{q}^{(n)}), \quad (1.46)$$

starting from any initial vector $\mathbf{q}^{(0)}$. The scheme will again converge to the limit $\lim_{n \rightarrow \infty} \mathbf{q}^{(n)} = \mathbf{q}$ when the matrix is stable. This iterative scheme shows that stability of the matrix $\mathbf{P} * \Gamma_S$ assures that the SDF in the valuation formula leads to sufficient discounting of future values, so that the valuation operator is sufficiently contractive. [Borovička and Stachurski \(2021\)](#) provide an exhaustive treatment of this problem.

Finally, we consider the case that allows cash flows to be non-stationary, with stationary growth rates

$$\log G_{t+1} - \log G_t = g_G(x_t, x_{t+1})$$

represented by a matrix Γ_G , with elements

$$[\Gamma_G]_{ij} = \exp(g_G(x_t = e_i, x_{t+1} = e_j)).$$

In this case, we need to rewrite the valuation formula (1.45) in price-dividend ratios by dividing it by G_t and manipulating the right-hand side to obtain

$$\frac{Q_t}{G_t} = E_t \left[\frac{S_{t+1}}{S_t} \left(\frac{Q_{t+1}}{G_{t+1}} + 1 \right) \frac{G_{t+1}}{G_t} \right]. \quad (1.47)$$

We conjecture that the price dividend ratio is stationary, and we can hence write

$$\frac{Q_t}{G_t} = q(x_t).$$

With the above-stated assumptions, the valuation formula then can be written as

$$q(x_t) = E_t [\exp(g_S(x_t, x_{t+1})) (q(x_{t+1}) + 1) \exp(g_G(x_t, x_{t+1}))]. \quad (1.48)$$

The argument in the expectation on the right-hand side is a function of x_t and x_{t+1} only, and the expectation integrates over x_{t+1} , which verifies that the left-hand side can be written only as a function of x_t , confirming the Markov structure for the price-dividend ratio.

The matrix implementation is analogous to the case with stationary cash flows. Denote

the \mathbf{q} the vector of state-dependent price-dividend ratios

$$[\mathbf{q}]_i = q(x_t = e_i).$$

Using the matrix structure, the recursive equation (1.47) can be expressed as

$$[\mathbf{q}]_i = q(x_t = e_i) = \sum_{j=1}^n [\mathbf{P}]_{ij} [\Gamma_S]_{ij} [\Gamma_G]_{ij} [\mathbf{q}]_j$$

This expression can be written in compact form as

$$\mathbf{q} = (\mathbf{P} * \Gamma_S * \Gamma_G) (\mathbf{q} + \mathbf{1}).$$

The equation has the solution

$$\mathbf{q} = [I - \mathbf{P} * \Gamma_S * \Gamma_G]^{-1} (\mathbf{P} * \Gamma_S * \Gamma_G) \mathbf{1} \quad (1.49)$$

provided that the matrix $\mathbf{P} * \Gamma_S * \Gamma_G$ is stable. The stability assumption requires that discounting in the stochastic discount factor Γ_S is sufficiently strong to overcome cash flow growth Γ_G . Alternatively, we could find the solution by iterating on the equation, analogously to (1.46).

1.5.2 Equity premium puzzle in the Mehra–Prescott economy

Mehra and Prescott (1985) construct an endowment economy with a 2-state Markov chain for the growth rate of aggregate endowment, which they equalize with aggregate consumption. They assume that investors have correct beliefs, $\mathbf{P} = \mathbf{P}^I$, with

$$\mathbf{P} = \begin{bmatrix} \phi & 1 - \phi \\ 1 - \phi & \phi \end{bmatrix} \quad (1.50)$$

and the growth rate of aggregate endowment is given by

$$\Gamma_G = \begin{bmatrix} 1 + \mu + \delta & 1 + \mu - \delta \\ 1 + \mu + \delta & 1 + \mu - \delta \end{bmatrix}.$$

The parameters of the two matrices are chosen to mimic, in a simple form, the dynamics of aggregate consumption in the U.S. economy, and they choose parameter values for an annual calibration of the model $\mu = 0.018$, $\delta = 0.036$, $\phi = 0.43$. The state e_1 is thus one with a high growth realization, while e_2 is the state with a low growth realization.

The investor is endowed with CRRA utility function with time preference β and risk aversion parameter γ . Mehra and Prescott (1985) explore asset pricing implications of this economy for a range of parameter values that, in their view, comfortably covers all plausible parametrizations, focusing on $\beta \in (0, 1)$ and $\gamma \in (0, 10)$.

Specifically, they analyze the equilibrium risk-free rate and risk premium on an asset that pays the aggregate consumption stream as its cash flow, $G_t = C_t$. The price of this

asset, Q_t , corresponds to the value all future consumption, which can be interpreted as total aggregate wealth. The price-dividend ratio $Q_t/G_t = Q_t/C_t$, which corresponds to the wealth-consumption ratio, is a function of the Markov state, $Q_t/C_t = q(x_t)$, as in expression (1.48), and represented as a vector \mathbf{q} computed using formula (1.49).

The risk-free rate is the return on the one-period risk-free bond

$$R_{t+1}^f = \left(E_t \left[\frac{S_{t+1}}{S_t} \right] \right)^{-1}$$

which can be expressed in matrix form as a matrix \mathbf{R}^f with elements

$$\mathbf{R}_{ij}^f = ([\mathbf{P} * \Gamma_S]_i \cdot \mathbf{1})^{-1}.$$

The fact that the return is safe is reflected in the fact that it only depends on the current state $x_t = e_i$, indicating that the return is known at time t .

To compute the average risk free rate, we average it across the stationary distribution π of the Markov chain. The stationary distribution is the eigenvector (normalized to make its elements sum up to one) of matrix \mathbf{P}' associated with the eigenvalue equal to one. Since \mathbf{P} defined in (1.50) is symmetric, we necessarily have $\pi = (0.5, 0.5)'$. We then have

$$E \left[R_{t+1}^f \right] = \sum_{i=1}^2 \pi_i ([\mathbf{P} * \Gamma_S]_i \cdot \mathbf{1})^{-1}.$$

The 'equity' return, or return on the claim on aggregate endowment, is given by

$$R_{t+1}^c = \frac{Q_{t+1} + C_{t+1}}{Q_t} = \frac{q(x_{t+1}) + 1}{q(x_t)} \frac{C_{t+1}}{C_t}.$$

In matrix notation,

$$\mathbf{R}_{ij}^c = \frac{\mathbf{q}_j + 1}{\mathbf{q}_i} [\Gamma_C]_{ij}.$$

The conditional equity premium is the difference between expectations of the two returns, conditional on a particular state today:

$$E_t \left[R_{t+1}^c - R_{t+1}^f | x_t = e_i \right] = \sum_{j=1}^2 \left(\mathbf{R}_{ij}^c - \mathbf{R}_{ij}^f \right) \mathbf{P}_{ij}$$

and the unconditional premium, under the unconditional distribution π ,

$$E \left[R_{t+1}^c - R_{t+1}^f | x_t = e_i \right] = \sum_{i=1}^2 \pi_i \sum_{j=1}^2 \left(\mathbf{R}_{ij}^c - \mathbf{R}_{ij}^f \right) \mathbf{P}_{ij}.$$

Mehra and Prescott (1985) then compute the average risk premium and average risk-free rate for a range of economies with preference parameters varied in the interval $\beta \in (0, 1)$, $\gamma \in (0, 10)$. The results are produced in Figure 1.1.

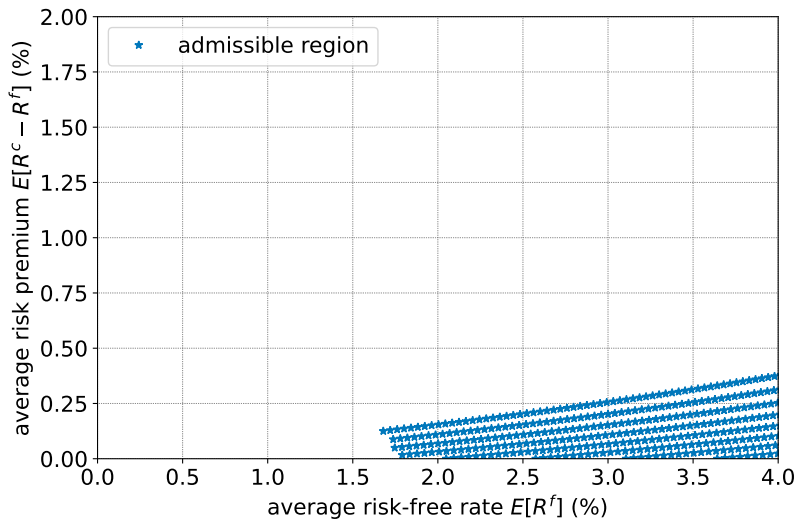


Figure 1.1: Combinations of the average risk-free rate and average risk premium obtained for a range of plausible preference parameters. Replication of Figure 4 in Mehra and Prescott (1985).

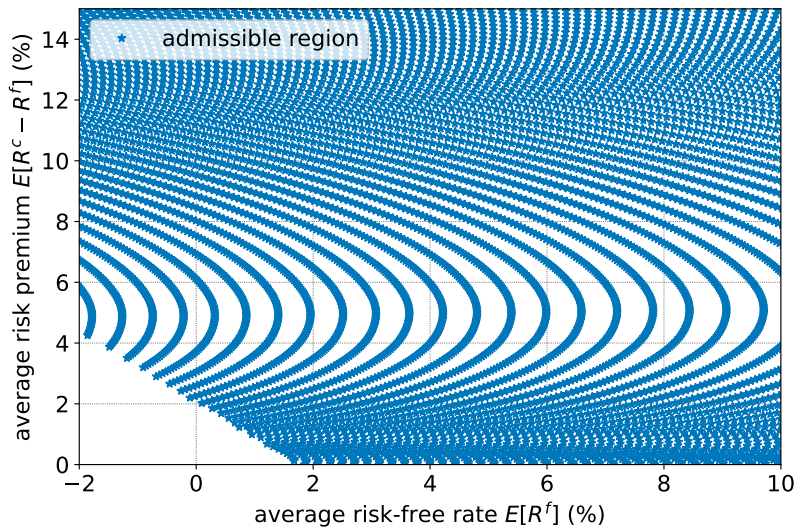


Figure 1.2: Replication of Figure 4 in Mehra and Prescott (1985) for a wider parameter range.

The blue region in the graph represents combinations of the average risk premium and average risk-free rate that can be obtained in such a range of economies. These are far away from empirical data, which put the average risk-free rate to about 1% and the average risk premium to about 6–7% annually. For Mehra and Prescott (1985), this is an equity premium puzzle, the apparent inability of a plausibly specified economy to replicate the low risk-free rate and high risk premium. Same conclusions have been reached under less restrictive modeling assumptions by Hansen and Singleton (1983).

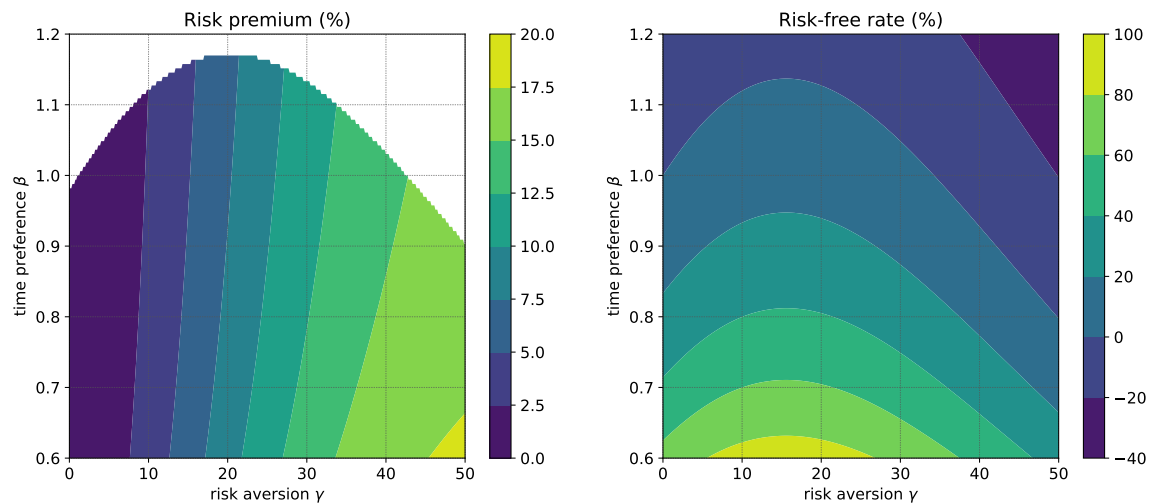


Figure 1.3: Risk premium and risk-free rate achieved for alternative parameterizations of the Mehra and Prescott (1985) economy. The white region in the left panel indicates parameter combinations for which the value of the claim on aggregate endowment is infinite.

One question is whether the model would fit empirical evidence if preferences parameters were allowed to be chosen in a broader range. Figure 1.2 replicates the exercise for the parameter range $\beta \in (0, 1.2)$ and $\gamma \in (0, 50)$. In this expanded parameter range, we can find an economy that does fit the two empirical moments.

Figure 1.3 indicates that an empirically sensible combination of the risk premium and risk-free rate is achieved for $\gamma \approx 20$ and $\beta \approx 1.12$. However, these parameters, with $\beta > 1$ and a very high risk aversion, would be considered outside the plausible range by essentially all researchers. More importantly, however, this economy at these preference parameter values will exhibit other features strongly at odds with data, like an extremely volatile risk-free rate. The conclusion from this exercise is that we need to resort to other models of preferences, sources of risk, or market arrangements, and the asset pricing literature has indeed made substantial progress in these directions.

1.6 Absence of arbitrage and equilibrium asset prices

In the endowment economy, we derived valuation formulas in the form of expected discounted values of cash flows for which the stochastic discount factor was associated with investors' marginal rate of substitution. However, the existence of valuation formulas

$$Q_t = E_t \left[\frac{S_{t+1}}{S_t} (G_{t+1} + Q_{t+1}) \right]$$

or, in returns form

$$1 = E_t \left[\frac{S_{t+1}}{S_t} R_{t+1} \right]$$

is not restricted to economies with markets populated by utility-maximizing investors. The existence of a strictly positive process S_t that satisfies these valuation restrictions is tightly related to absence of arbitrage in the given economy, without specifying an economic model that would imply the particular SDF.

In the endowment economy, we postulated a model of the SDF which, together with equilibrium restrictions, implied the behavior of asset prices. We now invert the problem. Imagine that we are given a set of asset price processes. What restrictions do these asset price processes need to satisfy so that they are arbitrage-free, in the sense that they do not permit investment strategies that would allow to generate positive profits out of a zero investment? It turns out that the existence of a strictly positive SDF is equivalent to absence of arbitrages. The existence of an SDF in the equilibrium economy is consistent with this result because arbitrages in such an economy would make the consumption-saving problem of the investor ill-posed.

The foundations of the link between the existence of a strictly positive stochastic discount factor and absence of arbitrage were developed in various setups by [Harrison and Kreps \(1979\)](#), [Harrison and Pliska \(1981\)](#), or [Kreps \(1981\)](#). We provide only a simple illustration of the results.

We restrict attention to a two-period market with K traded securities. There is a single state at time t and n possible states x_{t+1} at time $t + 1$. The probability distribution over states x_{t+1} is given by an $1 \times n$ vector $\mathbf{p} = (p_1, \dots, p_n)$. The securities $k = 1, \dots, K$ pay cash flows $G_{t+1}^k = G^k(x_{t+1})$ at time $t + 1$, given by $1 \times n$ vectors. We stack these K row vectors into a $K \times n$ matrix \mathbf{G} . The element \mathbf{G}_{kj} of this matrix represents the time- $t + 1$ payoff of asset k in state j . The time- t prices of these cash flows are denoted Q_t^k , stacked into a $K \times 1$ vector \mathbf{Q} .

At time t , the investor can choose a portfolio θ of the securities, represented by a $1 \times K$ vector, in which the element θ_k represents the number of units of asset k purchased. The time- t cost of such a portfolio is

$$\sum_{k=1}^K \theta_k Q_t^k = \sum_{k=1}^K \theta_k \mathbf{Q}_k = \theta \mathbf{Q},$$

and the payoff that this portfolio generates at time $t + 1$ is

$$\sum_{k=1}^K \theta_k G^k(x_{t+1}) = \theta \mathbf{G}.$$

We now formalize the notion of an arbitrage and a stochastic discount factor.

Definition 1.9. *An arbitrage is a portfolio θ such that either*

1. $\theta \mathbf{G} \in \mathbb{R}_+^n \setminus \{0\}$ and $\theta \mathbf{Q} \leq 0$; or
2. $\theta \mathbf{G} \in \mathbb{R}_+^n$ and $\theta \mathbf{Q} < 0$.

In words, an arbitrage is a portfolio that allows to generate a nonnegative payoff $\theta \mathbf{G}$

with strictly positive payoffs in some states out of a nonpositive investment $\theta \mathbf{Q}$ (condition 1), or a nonnegative payoff out of a strictly negative investment (condition 2).

Definition 1.10. A stochastic discount factor $s_{t+1} = s(x_{t+1})$, represented by an $1 \times n$ vector \mathbf{s} , is a strictly positive random variable such that

$$\mathbf{Q}_k = \sum_{j=1}^n \mathbf{p}_j \mathbf{s}_j \mathbf{G}_{kj} \quad k = 1, \dots, K. \quad (1.51)$$

The stochastic discount factor is thus a random variable which serves as a vector of state prices, such that the prices of all assets can be expressed as expected discounted payoffs, discounted by this SDF. In other words, an SDF can be interpreted as representing marginal valuations of payoffs in alternative states $j = 1, \dots, N$, without specific reference to a model that generates such marginal valuations.

Theorem 1.8 (Fundamental theorem of asset pricing). *A stochastic discount factor exists if and only if there is no arbitrage.*

The proof of this theorem is based on the so-called separating hyperplane theorem and can be found, for example, in [Duffie \(2001\)](#), Chapter 1.

The valuation formulas (1.51) are consistency conditions that state that prices of more complicated payoffs G_{t+1}^k can be represented as linear aggregates of prices of payoffs in individual states, given by the SDF \mathbf{s}_j . When there is no arbitrage, prices \mathbf{Q} must be mutually consistent in the sense that they allow such representation.

The theorem does not state whether such an SDF that supports prices of assets in a no-arbitrage market is unique. Uniqueness of the SDF is related to market completeness. When markets are complete, in the sense that any time $t + 1$ payoff can be replicated as a linear combination of payoffs of traded assets, then the SDF is unique. In the equilibrium economy studied in Section 1.3, we derived an SDF that is equalized with the marginal rate of substitution of the representative agent, and at least one such SDF must exist, otherwise the economy would exhibit arbitrages and hence could not be in equilibrium. However, there may be other stochastic processes that can constitute a valid SDF that delivers the same pricing implications of traded assets when markets are incomplete.

Chapter 2

Value function iteration in search problems

Textbook: Ljungqvist and Sargent (2018), Chapter 6.1–6.3 (p. 159–174).

Applications: McCall (1970).

QuantEcon: Quantitative Economics with Python, lectures 22–25.

Search theory is a fruitful area for applying the dynamic programming tools we study in this course. In this chapter, we formulate a simple model of dynamic choice, originally developed by McCall (1970), and illustrate how to approach the problem recursively. This model is the foundation of a large literature on the so-called Diamond–Mortensen–Pissarides search-and-matching theory of the labor market. We want to understand the following.

- How to approach the problem in a smart, tractable way?
- What can we say about the existence of solutions to this problem? Under what conditions?
- Can we characterize the behavior of this model on theoretical grounds?
- What are useful computational methods to attack this problem?

In this chapter, we focus on providing economic intuition, before delving into a formal analysis of dynamic programming in following chapters.

2.1 A simple model of intertemporal search

In the McCall (1970) model, an unemployed worker is searching for a job. Every period $t = 0, 1, 2, \dots$, she receives a job offer, for which a wage is drawn from a distribution with cdf $F(w)$, with $F(0) = 0$, $F(B) = 1$ for some $B > 0$. These offers are independently

and identically distributed (iid) over time. If the distribution has a density, we denote it $f(w) = F'(w)$ but the existence of a density is not required for most of the results.

When the worker accepts an offer, she stays at the job forever at the constant wage w . When the worker rejects, she continues to search next period, when a new offer is drawn. The worker maximizes the present discounted value of future income

$$E \left[\sum_{t=0}^{\infty} \beta^t y_t \right], \quad \beta \in (0,1)$$

where $y_t = w$ when employed, and $y_t = c \geq 0$ when unemployed. The parameter c can be interpreted as a flow of unemployment benefits. The parameter β represents the rate of **geometric discounting**.

The expectations operator $E[\cdot]$ represents the mathematical expectation of a random variable. In particular, the expected value of the wage w before it is drawn is

$$E[w] = \int_0^B w dF(w) = \int_0^B w f(w) dw.$$

The goal is to devise the optimal strategy for which offers to accept and when. More formally, we are solving the **sequence problem**

$$V_0^* = \max_{\{a_t\}_{t=0}^{\infty}} E \left[\sum_{t=0}^{\infty} \beta^t y_t \right] \quad (2.1)$$

where $a_t \in \{\text{accept, reject}\}$ if the worker has not yet accepted any other offer before time t , and $a_t \in \{\}$ otherwise. V_0^* is the **value function** associated with the worker's problem, i.e., the present discounted value of wages attained if the worker behaves optimally. Here, we will make the timing assumption that V_0^* is evaluated after the worker learned the first wage offer w_0 but before she made the decision whether to accept or reject it. V_0^* hence depends on w_0 .

2.1.1 Information sets, beliefs, and strategies

Notice that the structure of the problem is complex. For every realized sequence of wage offers, there can be at most one offer that is accepted, since after the offer is accepted, the worker stays in the job forever. Now consider a worker who is still unemployed at time t and who receives an offer w_t . Her decision whether to accept or not (**acceptance decision**) can now depend on anything she observed so far:

- the current wage offer w_t ,
- the history of previously rejected offers $(w_0, w_1, \dots, w_{t-1}) \doteq w^{t-1}$,
- time t ,
- potentially other observed information.

These pieces of information constitute the worker's **information set** at time t . Notice that as time passes, the worker adds more and more information into this information set. We will make the (strong) assumption that the worker has **perfect memory** and never forgets any of the information.

The acceptance decision also depends on worker's **beliefs** about future offers. Above, we have described the offers as iid draws from a distribution $F(w)$. In this problem, we can think about these offers as being drawn by **nature**. The sequence of these offers then forms a stochastic process called the **data-generating process** with a probability distribution determined by nature.

By specifying the value function (2.1) using the mathematical expectation $E[\cdot]$, we implicitly assumed that the worker has **correct beliefs** that are aligned with the data-generating process. This means that $F(\bar{w})$ represents both the objective probability that the drawn wage offer will satisfy $w \leq \bar{w}$, as well as the worker's belief that $w \leq \bar{w}$. Imposing that the two probability measures are the same is a potentially strong (but often convenient and informative) assumption that lies at the heart of **rational expectations** models.

The period- t income y_t depends on which offer has been accepted and when. The acceptance decision is the only relevant decision to be made in this problem. A complete description of worker's decisions in every contingency is called a **strategy**.

2.1.2 Recursive formulation

We now show how to simplify the problem by first deciding which information is really relevant for the worker's decision, and then relying on a recursive formulation. The recursive approach splits the problem into a decision today and a continuation problem for the next period.

Denote $E_t[\cdot]$ the expectations operator that conditions on all information the worker has available at time t . Intuitively,

$$\begin{aligned} V_0^* &= \max_{\{a_t\}_{t=0}^{\infty}} \left\{ y_0 + \beta E_0 \left[\sum_{t=1}^{\infty} \beta^{t-1} y_t \right] \right\} = \max_{a_0} \left\{ y_0 + \beta \max_{\{a_t\}_{t=1}^{\infty}} E_0 \left[\sum_{t=1}^{\infty} \beta^{t-1} y_t \right] \right\} \quad (2.2) \\ &= \max_{a_0} \left\{ y_0 + \beta E_0 \left[\max_{\{a_t\}_{t=1}^{\infty}} \left\{ y_1 + \beta E_1 \left[\sum_{t=2}^{\infty} \beta^{t-2} y_t \right] \right\} \right] \right\} = \max_{a_0} \{ y_0 + \beta E_0 [V_1^*] \}. \end{aligned}$$

Observe that V_0^* conditions on information available at time 0, while V_1^* conditions on all information available up to time t . In order to make this problem recursive, we need to make sure that we find representation in which the information that the worker uses to make a decision has the same structure each period.

First consider the relevant information for the worker's decision. We conjecture that when an unemployed worker receives an offer w in period t , the only relevant information for whether to accept or reject this offer is w . In other words, w summarizes everything that describes worker's current situation, and that is why we call w the **state** (or state variable). **Finding the state is an art**, and judicious choice of the state can dramatically

simplify many interesting problems.

Is this guess justified? Consider what other information could the decision depend on:

- Past offers w^{t-1} . These would be relevant in two cases. First, if the agent could revisit and accept those old offers at time t . That would be a model with offer recall. Second, if those old offers helped the agent predict the distribution of future offers. That would be a model with learning. In both cases, we would need to devise a way how to encode these past offers w^{t-1} as a part of the state at time t .
- Time t . That would be relevant, for instance, in a situation when the agent faces a finite-horizon problem, e.g., in the case of a finite lifetime. Another example is a model with seasonality, where the distribution of offers depends on the calendar index of the month or quarter. However, our problem is **time-invariant** and hence time does not show up as a relevant state.

With this conjecture for the state, we can denote $V(w)$ the present discounted value of income for an unemployed worker who currently has an offer w and decides whether to accept or reject.

If the worker with wage offer w accepts, she becomes employed and receives w forever. We denote $V^a(w)$ the present discounted value of these wages:

$$V^a(w) = \sum_{t=0}^{\infty} \beta^t w = \frac{w}{1-\beta}.$$

When the worker rejects, she receives c , and the next period she draws a new random offer w' . The value next period from this random offer is $V(w')$. Hence in the current period, the value of rejecting the current offer w is

$$c + \beta \int_0^B V(w') dF(w').$$

The worker can choose between accepting and rejecting, and hence

$$V(w) = \max_{\{\text{accept, reject}\}} \left\{ V^a(w), c + \beta \int_0^B V(w') dF(w') \right\}. \quad (2.3)$$

Observe that the structure of the decision problem satisfies our conjecture that all that matters for the current decision is the current wage offer w . Hence indeed w is the state for the problem.

Further notice that the problem is now written **recursively**. Optimal strategy that generates the value $V(w)$ consists of a decision today and a **continuation problem** next period. The decision problem simplified significantly because it now consists of a single accept/reject decision. The additional complication comes now from the fact that we have the unknown function $V(w)$ both on the left and the right of equation (2.3). This **functional equation** is called the **Bellman equation**. It is a functional equation because its solution is the whole function $V(w)$. Associated with the solution is the **optimal decision**

rule (or policy function) — a function $a(w) \in \{\text{accept, reject}\}$ that states whether to accept or reject the wage offer, for every feasible offer w . This is in contrast to the sequence problem (2.2) where the solution $\{V_t^*\}_{t=0}^\infty$ and optimal decisions $\{a_t^*\}_{t=0}^\infty$ are stochastic processes that depend on the history of wage draws w^t .

Equation (2.3) can be rewritten by defining an operator T that takes as input a function $v(w)$ defined on $[0, B]$ and produces a new function Tv defined as:

$$(Tv)(w) = \max_{\{\text{accept, reject}\}} \left\{ \frac{w}{1-\beta}, c + \beta \int_0^B v(w') dF(w') \right\}.$$

This operator is called the **Bellman operator**, and solving (2.3) corresponds to finding the fixed point that satisfies

$$V = TV.$$

Notice that our ultimate goal is to find the value function V_0^* in (2.2) for every given w_0 . So we will need to show that the problem (2.3) is equivalent to (2.2), and that $V(w_0) = V_0^*$ when the time-0 offer is w_0 . Similarly, we want to show that the policy function $a(w_0)$ is the same as the time-0 decision a_0^* when the offer is w_0 .

2.1.3 Principle of optimality

The recursive representation (2.3) is the foundation of the **dynamic programming** method and the validity of this approach is based on the **principle of optimality**. This principle, due to Richard Bellman, breaks down the infinite-dimensional problem (2.1) into smaller subproblems:

An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision. (Bellman (1957), Chapter III.3)

In (2.3), the value $V(w')$ on the right-hand side is the value of the optimal policy conditional on drawing w' next period, where optimal decisions are followed in the next and all subsequent periods. This means that in order to determine the optimal decision in the current period, one needs to determine its consequences for the **state** next period and then work with **next period value** $V(w')$ computed under an optimized subsequent path, given the next period state w' . In this example, the state w' is independent of the current period action but this is not necessarily true in general.

This is a massive computational simplification but it does not come for free. First, the recursive formulation crucially relies on the ability to determine the **state variable**. Second, equation (2.3) may still not be easy to solve because it involves computing a function that is a **fixed point** of that equation. Finally, as discussed in the previous section, we will need to verify that this fixed point $V(w)$ indeed corresponds to the value function V_0^* . Theoretical results that establish existence and uniqueness of such fixed points and verification theorems that show equivalence between the fixed point and the value function often require imposing assumptions that are quite restrictive for economic applications.




Figure 2.1: Acceptance decision in the McCall (1970) model.

2.1.4 Characterizing optimal policy

In order to determine the acceptance decision, consider Figure 2.1. The solid black line plots the value of accepting the offer $V^a(w) = w / (1 - \beta)$ as a function of w . This is the first term on the right-hand side of (2.3). The dashed line plots the value of rejecting the offer w . This is the second term on the right-hand side of (2.3), which we denote

$$Q = c + \beta \int_0^B V(w') dF(w'). \quad (2.4)$$

Notice that at this point, Q is unknown but we know it is independent of w .

The value $V(w)$ is then represented by the maximum of the two functions, depicted in red. Since $V^a(w)$ is linear and increasing while Q is constant, the decision rule takes the form of a **reservation wage** \bar{w} such that the agent accepts the offer if $w > \bar{w}$ and rejects it if $w < \bar{w}$. When $w = \bar{w}$, the agent is indifferent between rejecting and accepting.

The function $V(w)$ is therefore given by

$$V(w) = \begin{cases} c + \beta \int_0^B V(w') dF(w') = \frac{\bar{w}}{1 - \beta} & \text{if } w \leq \bar{w} \\ \frac{w}{1 - \beta} & \text{if } w \geq \bar{w} \end{cases} \quad (2.5)$$

where \bar{w} is unknown.

While the form of the decision rule is clear, it could in principle be the case that there are two different value functions V_1 and V_2 consistent with alternative reservation wages \bar{w}_1 and \bar{w}_2 that solve the functional equation (2.5). To prove uniqueness of the solution, we convert the functional equation for $V(w)$ into a simple equation for \bar{w} . The first line in




Figure 2.2: Characterizing the reservation wage in the McCall (1970) model.

(2.5) implies

$$\begin{aligned} \frac{\bar{w}}{1-\beta} &= c + \beta \int_0^{\bar{w}} \frac{\bar{w}}{1-\beta} dF(w') + \beta \int_{\bar{w}}^B \frac{w'}{1-\beta} dF(w') \\ &= c + \beta \int_0^B \frac{\bar{w}}{1-\beta} dF(w') + \frac{\beta}{1-\beta} \int_{\bar{w}}^B (w' - \bar{w}) dF(w') \end{aligned}$$

and hence

$$\bar{w} - c = \frac{\beta}{1-\beta} \int_{\bar{w}}^B (w' - \bar{w}) dF(w'). \quad (2.6)$$

The only unknown in this equation is the reservation wage \bar{w} . Observe that the left-hand side represents the cost of searching one more time when the current wage offer is \bar{w} . It consists of sacrificing \bar{w} net of the unemployment benefit c received when rejecting the offer. The right-hand side is the net benefit from searching one more time, consisting of the present discounted value of potential wages in excess of \bar{w} . Equality between costs and benefits is the usual optimality condition on the choice of \bar{w} .

We plot the two sides of equation (2.6) in Figure 2.2. The left-hand side is the linear increasing function $w - c$, while the right-hand side is given by

$$h(w) = \frac{\beta}{1-\beta} \int_w^B (w' - w) dF(w'). \quad (2.7)$$

The reservation \bar{w} is the value of w where these two functions intersect. Obviously, $h(0) = \frac{\beta}{1-\beta} E[w']$ and $h(B) = 0$, and applying the Leibniz rule, we get

$$\begin{aligned} h'(w) &= -\frac{\beta}{1-\beta} \int_w^B dF(w') = -\frac{\beta}{1-\beta} [1 - F(w)] < 0 \\ h''(w) &= \frac{\beta}{1-\beta} F'(w) > 0 \end{aligned}$$

hence the function $h(w)$ is decreasing and convex. Therefore, as long as $0 < c < B$ and the distribution of w has full support on $[0, B]$, there is a unique interior reservation wage \bar{w} . The potential existence of mass points in the distribution $F(w)$ does not alter this result but if the solution of equation (2.6) lies in an interval $[w_l, w_h]$ on which $F(w)$ has zero mass, then any $w \in [w_l, w_h]$ is a valid reservation wage.

Introspection of equation (2.6) also yields the following associated results:

- $\bar{w} > c$ as long as $F(c) < 1$. This means that there is **option value of waiting** and the worker rejects some offers higher than c in order to wait for better future offers.
- Equation (2.6) implies that \bar{w} does not depend on the shape of $F(w)$ on $[0, \bar{w})$ because the right-hand side is not integrating over that interval. Economically, the cutoff does not depend on offers that get rejected anyway.
- $d\bar{w}/dc > 0$. Observe that only the increasing curve $w - c$ is a function of c . An increase in c shifts the curve down, which shifts \bar{w} to the right. A higher unemployment benefit makes workers pickier (and consequently also increases length of unemployment).
- $d\bar{w}/d\beta > 0$. An increase in β rescales the decreasing function $h(w)$, making it steeper. Consequently, \bar{w} shifts to the right. Higher patience increases the **option value of waiting**.

2.1.5 Mean-preserving spreads

Recall that the draw of an offer in the McCall (1970) model is a random variable w with distribution $F(w)$ defined as $P(w \leq \tilde{w}) = F(\tilde{w})$, with support on $[0, B]$. The mean is defined as $E[w] = \int_0^B w dF(w)$. Integration by parts then yields

$$\begin{aligned} E[w] &= \int_0^B w dF(w) = wF(w) \Big|_0^B - \int_0^B F(w) dw = \\ &= B - \int_0^B F(w) dw = \int_0^B [1 - F(w)] dw. \end{aligned} \quad (2.8)$$

Consider now two independent draws w_1, w_2 . Then

$$\begin{aligned} P(\max(w_1, w_2) \leq w) &= P((w_1 \leq w) \cap (w_2 \leq w)) = \\ &= P(w_1 \leq w) P(w_2 \leq w) = F(w)^2. \end{aligned}$$

Consequently, (2.8) implies that

$$E[\max(w_1, w_2)] = B - \int_0^B F(w)^2 dw \quad (2.9)$$

or, more generally, the expected value of the maximum of n draws,

$$M_n \doteq E [\max (w_1, \dots, w_n)] = B - \int_0^B F(w)^n dw.$$

Rothschild and Stiglitz (1970) introduced the concept of a **mean-preserving spread** as a way of describing whether one random variable is more ‘variable’ than another. Consider two random variables X and Y . Rothschild and Stiglitz (1970) have shown that the following concepts are equivalent:

- a) $E [u (X)] \geq E [u (Y)]$ for all concave u , i.e., all risk averters prefer X to Y .
- b) $Y = X + Z$ where $E [Z | X] = 0$.
- c) If X and Y have a compact support $[a, b]$ and distribution functions F_X and F_Y , respectively, then

$$T(\tilde{z}) = \int_a^{\tilde{z}} [F_Y(z) - F_X(z)] dz$$

satisfies $T(z) \geq 0, \forall z \in [a, b]$ and $T(b) = 0$, i.e., Y has more weight in tails than X but the same mean as X .

In this case, we say that Y was constructed from X by means of a mean-preserving spread. Intuitively, a mean-preserving spread makes the distribution of the random variable more uncertain, while preserving its mean. In the definitions above, property b) states that Y can be constructed from X by perturbing X with a random noise Z whose mean cannot be predicted by X (but higher moments potentially can). Property a) states that all risk averse agents dislike Y relative to X .

In particular, the concept of a mean preserving is broader than a pure comparison of variances of X and Y . We now utilize this concept, in particular property c) above, to study implications of an increase in uncertainty in the wage-offer distributions in the context of the McCall (1970) model.

Consider random variables w_1 and w_2 with distribution functions F_1 and F_2 . First, observe that (2.8) implies that w_1 and w_2 having the same mean is equivalent to

$$0 = E [w_1] - E [w_2] = \int_0^B [F_2(w) - F_1(w)] dw, \quad (2.10)$$

so indeed $T(B) = 0$. Second, the two distributions F_1 and F_2 are said to satisfy the **single-crossing property** if there exists \hat{w} such that

$$F_2(w) - F_1(w) \begin{cases} \geq 0 & \text{when } w \leq \hat{w} \\ \leq 0 & \text{when } w \geq \hat{w} \end{cases} \quad (2.11)$$

(or vice versa). Graphically, the single-crossing property is depicted in Figure 2.3. Since a high slope of the distribution function indicates parts of the support that have high density




Figure 2.3: Two distributions F_1 and F_2 satisfying the single-crossing property.

of the probability distribution, the graph shows that w_2 indeed has ‘more weight in tails’ as property c) indicates.

The two conditions (2.10) and (2.11) imply that

$$\int_0^{\tilde{w}} [F_2(w) - F_1(w)] dw \geq 0 \quad 0 \leq \tilde{w} \leq B. \quad (2.12)$$

Observe that condition (2.11) implies that the integral in (2.12) has to first increase and then decrease. Condition (2.10) states that the integral has to keep decreasing until it reaches zero at $\tilde{w} = B$. Rothschild and Stiglitz (1970) then use conditions (2.10) and (2.12) to define the **mean-preserving spread**: w_2 has been obtained from w_1 by a mean preserving spread if the two distributions satisfy (2.10) and (2.12).¹

Diamond and Stiglitz (1974) derive a differential measure of the mean-preserving spread. In order to do so, consider a family of distributions $F(w; r)$ indexed by r . Assume that for $r_2 > r_1$, the two distributions satisfy conditions (2.10) and (2.12), i.e.,

$$\begin{aligned} \int_0^B [F(w; r_2) - F(w; r_1)] dw &= 0 \\ \int_0^{\tilde{w}} [F(w; r_2) - F(w; r_1)] dw &\geq 0 \quad 0 \leq \tilde{w} \leq B. \end{aligned}$$

¹The preference for using conditions (2.10) and (2.12) rather than (2.10) and (2.11) lies in the fact that (2.10) and (2.12) satisfy transitivity. Specifically, if w_2 was obtained from w_1 by a mean-preserving spread, and w_3 from w_2 by another mean-preserving spread, then w_3 and w_1 do not need to satisfy the single-crossing property but they will satisfy (2.10) and (2.12), so these properties can be used to order distributions.

Then a differential version of these conditions reads

$$\begin{aligned} \int_0^B F_r(w; r) dw &= 0 \\ \int_0^{\bar{w}} F_r(w; r) dw &\geq 0 \quad 0 \leq \bar{w} \leq B. \end{aligned}$$

2.1.6 Effects of mean-preserving spread in the McCall (1970) model

Let us return to equation (2.6) representing the optimal choice of the reservation wage \bar{w} :

$$\bar{w} - c = \frac{\beta}{1 - \beta} \int_{\bar{w}}^B (w' - \bar{w}) dF(w'),$$

and rewrite the right-hand side using integration by parts as

$$\begin{aligned} \bar{w} - c &= \frac{\beta}{1 - \beta} (E[w'] - \bar{w}) - \frac{\beta}{1 - \beta} \int_0^{\bar{w}} (w' - \bar{w}) dF(w') \\ &= \frac{\beta}{1 - \beta} (E[w'] - \bar{w}) - \frac{\beta}{1 - \beta} \underbrace{(w' - \bar{w}) F(w') \Big|_0^{\bar{w}}}_{=0} + \frac{\beta}{1 - \beta} \int_0^{\bar{w}} F(w') dw' \end{aligned}$$

or

$$\bar{w} - c = \beta (E[w'] - c) + \beta \int_0^{\bar{w}} F(w') dw'. \quad (2.13)$$

Defining

$$g(w) = \beta (E[w'] - c) + \beta \int_0^w F(w') dw'$$

we plot in Figure 2.4 the two sides of equation (2.13), $w - c$ and $g(w)$. Notice that $g(0) = \beta (E[w'] - c)$ could be negative or positive but certainly higher than $-c$. We know that these two curves cross exactly once since we already determined there is a unique solution \bar{w} .

A mean-preserving spread from $F(w)$ to $\tilde{F}(w)$ increases the last term in $g(w)$. The new curve is depicted as $\tilde{g}(w)$ in Figure 2.4. Clearly, $g(0) = \tilde{g}(0)$ and, from the definition of the mean-preserving spread, $g(B) = \tilde{g}(B)$. As a result, the reservation wage shifts to the right to $\tilde{\bar{w}}$.

Economically, a more dispersed distribution of wage offers increases the option value of waiting for a better offer to arrive. The fact that bad offers got worse than before is irrelevant because those offers would have been rejected anyway. It is the upside from the improvement in the good offers that matters for the decision.

2.1.7 Welfare

A natural question is to ask what are the implications of the changes in the model for the welfare of the worker. Here, $V(w)$ measures the present discounted value of wages and




Figure 2.4: Characterizing the role of mean-preserving spread in the [McCall \(1970\)](#) model. A mean-preserving spread in the distribution of wage offers increases the reservation wage.

unemployment benefits under the optimal choice of the worker, for a worker with offer w at hand. When these wages and benefits are measured in units of real consumption, it is an appropriate measure of welfare for a risk-neutral worker.

Clearly, an increase in unemployment benefits c increases welfare of the worker. The present discounted value of consumption increases with an increase in c even if the worker leaves her policy rule unchanged. A reoptimization of the acceptance decision can then increase welfare even further.

What about an increase in mean-preserving spread? Recall that the value function is written down in (2.5):

$$V(w) = \begin{cases} c + \beta \int_0^B V(w') dF(w') = \frac{\bar{w}}{1-\beta} & \text{if } w \leq \bar{w} \\ \frac{w}{1-\beta} & \text{if } w \geq \bar{w} \end{cases}$$

Since \bar{w} increased under the mean-preserving spread, $V(w)$ must (nonstrictly) increase as well. Again, this is the effect of the improved option value of waiting.

However, notice that under the mean-preserving spread, the distribution of wage offers also changed. Hence it is relevant to ask whether agent's welfare increases even before she knows the wage offer in the given period. The relevant object to study therefore is

$$\int_0^B V(w) dF(w)$$

Manipulating this expression, we yield

$$\begin{aligned}
 \int_0^B V(w) dF(w) &= \int_0^{\bar{w}} \frac{\bar{w}}{1-\beta} dF(w) + \int_{\bar{w}}^B \frac{w}{1-\beta} dF(w) = \\
 &= \frac{\bar{w}}{1-\beta} F(\bar{w}) + \frac{w}{1-\beta} F(w) \Big|_{\bar{w}}^B - \frac{1}{1-\beta} \int_{\bar{w}}^B F(w) dw \\
 &= \frac{B}{1-\beta} - \frac{1}{1-\beta} \int_0^B F(w) dw + \frac{1}{1-\beta} \int_0^{\bar{w}} F(w) dw
 \end{aligned}$$

where the first equality uses the fact that $V(w)$ is a piecewise linear function, the second equality uses integration by parts, and the third rewrites the last integral from the second line.

Now recall from (2.10) and (2.12) that a mean preserving spread leaves the second term on the last line unchanged and, holding the upper bound \bar{w} fixed, increases the last term. At the same time, the fact that the reservation wage \bar{w} also increases under the optimal policy, further increases the last term. So even *ex ante*, before the worker receives the offer in the given period, her welfare increases.

In fact, there is also a very simple answer to this question. Since we have established that the reservation wage \bar{w} increases after the mean reserving spread, and, from (2.5),

$$c + \beta \int_0^B V(w') dF(w') = \frac{\bar{w}}{1-\beta}$$

it must be that $\int_0^B V(w') dF(w')$ also increases.

2.2 Extensions

Here, we consider some extensions that document the power of the recursive approach, and, at the same time, highlight some challenges.

2.2.1 Option to quit

Consider the same problem as before but now the worker has the option to quit the existing job at the end of each period, stay unemployed for one period and then start searching again as an unemployed worker at the beginning of the subsequent period. Once she decides to quit, she has no option to come back to her old job.

Recall that we have established that the optimal reservation wage \bar{w} satisfies

$$V(\bar{w}) = \frac{\bar{w}}{1-\beta} = c + \beta \int_0^B V(w') dF(w').$$

It is clear that with this new option, the optimal policy still has the form of a reservation wage \bar{w} (potentially different than before). Given a wage offer w , the agent has now three

options:

- Accept the wage and keep the job forever, yielding

$$\frac{w}{1 - \beta}$$

- Reject the wage, yielding

$$c + \beta \int_0^B V(w') dF(w') = \frac{\bar{w}}{1 - \beta}$$

- Accept the wage but quit after t periods, then follow the optimal policy. This yields

$$\begin{aligned} \frac{1 - \beta^t}{1 - \beta} w + \beta^t \left(c + \beta \int_0^B V(w') dF(w') \right) &= \frac{1 - \beta^t}{1 - \beta} w + \beta^t \frac{\bar{w}}{1 - \beta} \\ &= \frac{w}{1 - \beta} - \beta^t \frac{w - \bar{w}}{1 - \beta} \end{aligned}$$

Observe that due to time-invariance, the value of having an offer w is $V(w)$ irrespective in which period the offer is made.

Hence quitting is never optimal because the value of accepting and quitting later is either lower than the value of accepting and never quitting (in the case when $w > \bar{w}$) or lower than rejecting and waiting ($w < \bar{w}$). The only point of indifference is $w = \bar{w}$ in which case all three policies yield the same value. Consequently, the option value of quitting is zero and the option will never be exercised.

The experiment we conducted is closely related to the **principle of optimality** we discussed in Section 2.1.3. We conjectured an optimal policy (accepting when $w \geq \bar{w}$ and never quitting) and then constructed deviations from this policy to confirm that these deviations do not lead to improvements.

2.2.2 Firing

Consider now the case when the worker can be fired. In particular, at the end of every period when the worker was employed, she can be fired with probability α , in which case she enters the next period as unemployed, collects unemployment benefits for one period, and then starts receiving offers from the following period onward. Denote $\hat{V}(w)$ the value for an unemployed worker with offer w at hand. We have the Bellman equation

$$\hat{V}(w) = \max_{\{\text{accept, reject}\}} \{w + \beta [(1 - \alpha) \hat{V}(w) + \alpha (c + \beta E[\hat{V}(w')])], c + \beta E[\hat{V}(w')]\} \quad (2.14)$$

where $E[\hat{V}(w')] = \int_0^B \hat{V}(w') dF(w')$. The fact that we can insert $\hat{V}(w)$ at the right-hand side is based on the observation that an offer w that is accepted this period would have

been accepted next period as well, so as long as the worker stays at the job, her value continues to be $\hat{V}(w)$.

Remark 2.1. Sometimes, it may be more transparent to split the Bellman equation (2.14) into components. Define $\hat{V}^e(w)$ the value of an employed worker who did not get fired at the beginning of the current period and hence will continue working at wage w . Then (2.14) can be rewritten as a pair of equations

$$\begin{aligned}\hat{V}(w) &= \max_{\{\text{accept, reject}\}} \{ \hat{V}^e(w), c + \beta E [\hat{V}(w')] \} \\ \hat{V}^e(w) &= w + \beta [(1 - \alpha) \hat{V}^e(w) + \alpha (c + \beta E [\hat{V}(w')])].\end{aligned}$$

This is particularly useful in situations when the state of the economy changes from period to period, and hence offers that are accepted this period may not be accepted in the next period.

Remark 2.2. If the fired worker did not have to go through one period of unemployment but could receive an offer right away, the Bellman equation (2.14) would be given by

$$\hat{V}(w) = \max_{\{\text{accept, reject}\}} \{ w + \beta [(1 - \alpha) \hat{V}(w) + \alpha E [\hat{V}(w')]] , c + \beta E [\hat{V}(w')] \}.$$

Again, the optimal policy is a reservation wage \bar{w} . This implies

$$\hat{V}(w) = \begin{cases} \frac{w + \beta \alpha (c + \beta E [\hat{V}(w')])}{1 - \beta (1 - \alpha)} & \text{if } w \geq \bar{w} \\ c + \beta E [\hat{V}(w')] & \text{if } w \leq \bar{w} \end{cases}$$

and the reservation wage must solve

$$\frac{\bar{w} + \beta \alpha (c + \beta E [\hat{V}(w')])}{1 - \beta (1 - \alpha)} = c + \beta E [\hat{V}(w')]$$

or, after rearranging,

$$\frac{\bar{w}}{1 - \beta} = c + \beta E [\hat{V}(w')] \quad (2.15)$$

This expression has a form that is identical to the reservation wage condition from the first line of (2.5) from the model without firing. However, since $\hat{V}(w)$ and $V(w)$ differ, so does the reservation wage. In particular, $\hat{V}(w) \leq V(w)$. This is easy to see, since we have shown that it is not optimal to quit — hence firing must be undesirable to the worker. In general, this inequality is strict (except pathological cases), and this is true even for wage offers w that get rejected, since the possibility of firing reduces the option value of future acceptable offers.

Since $\hat{V}(w) \leq V(w)$, then also the reservation wage \bar{w} is lower when firing is possible. This follows directly from (2.15). The benefit of accepting a job offer comes from the surplus on top of staying unemployed. Probability of firing decreases the most the surplus for the jobs with highest wages. Hence if the best jobs are less attractive, there is less of a reason to wait for those jobs, and the worker is willing to settle for a worse wage offer.

2.2.3 Finite horizon

The previous examples assumed that the time horizon is infinite, which lead to time-invariant optimal policy and associated function $V(w)$. However, many problems, for example life-cycle problems, are problems with finite horizons. Consider a problem of a worker who lives and works in periods $t = 0, \dots, T$ and retires after T . Then the state is (w, t) and optimal policy is an explicit function of time. In particular,

$$\begin{aligned} V_t(w) &= \max_{\{\text{accept, reject}\}} \left\{ \frac{1 - \beta^{T-t+1}}{1 - \beta} w, c + \beta \int_0^B V_{t+1}(w') dF(w') \right\} \quad t = 0, \dots, T \\ V_{T+1}(w) &= 0 \end{aligned} \quad (2.16)$$

Given the structure of the problem, the optimal policy is given by a sequence of reservation wages \bar{w}_t and associated values $V_t(w)$, $t = 0, \dots, T$. The reservation wage \bar{w}_t at time t must satisfy

$$\frac{1 - \beta^{T-t+1}}{1 - \beta} \bar{w}_t = c + \beta \int_0^B V_{t+1}(w') dF(w').$$

Multiplying by $1 - \beta$ and subtracting $(1 - \beta^{T-t+1})c$, we get

$$\begin{aligned} (1 - \beta^{T-t+1})(\bar{w}_t - c) &= (1 - \beta)c - (1 - \beta^{T-t+1})c + \beta(1 - \beta) \int_0^B V_{t+1}(w') dF(w') \\ &= -\beta(1 - \beta^{T-t})c + \beta(1 - \beta) \int_0^{\bar{w}_{t+1}} \frac{1 - \beta^{T-t}}{1 - \beta} \bar{w}_{t+1} dF(w') \\ &\quad + \beta(1 - \beta) \int_{\bar{w}_{t+1}}^B \frac{1 - \beta^{T-t}}{1 - \beta} w' dF(w') \\ &= \beta(1 - \beta^{T-t}) \left(E[w'] - c - \int_0^{\bar{w}_{t+1}} (w' - \bar{w}_{t+1}) dF(w') \right). \end{aligned}$$

Applying integration by parts to the integral and dividing by $1 - \beta^{T-t+1}$, we obtain

$$\bar{w}_t - c = \beta \frac{1 - \beta^{T-t}}{1 - \beta^{T-t+1}} \left(E[w'] - c + \int_0^{\bar{w}_{t+1}} F(w') dw' \right). \quad (2.17)$$

We can now compare this expression to the infinite-horizon formula (2.13). These expressions are almost identical, except the term

$$\frac{1 - \beta^{T-t}}{1 - \beta^{T-t+1}} < 1.$$

We can now characterize \bar{w}_t .

- Obviously, $\bar{w}_T = c$. Any offer better than the unemployment benefit will be accepted in the last period, since there is no option value of waiting.
- Next, as $T \rightarrow \infty$, we have $\bar{w}_t \rightarrow \bar{w}$ for every fixed t . This follows from taking the limit in (2.17). When the terminal period is far in the future, the reservation wage

approaches the time-invariant solution.

- Finally, the reservation wage is decreasing in time t (i.e., decreasing as the worker approaches the terminal period). To show this, we proceed as follows. Denote

$$\tilde{\beta}_t = \beta \frac{1 - \beta^{T-t}}{1 - \beta^{T-t+1}}$$

For every t , equation (2.17) has a unique fixed point \tilde{w}_t that satisfies

$$\tilde{w}_t - c = \tilde{\beta}_t \left(E[w'] - c + \int_0^{\tilde{w}_t} F(w') dw' \right).$$

This is easy to see since it corresponds to equation (2.13) with an augmented time-preference coefficient $\tilde{\beta}_t$. In particular, this fixed point is always larger than c . Next, denote the right-hand side of (2.17) plus c as

$$g_t(\bar{w}_{t+1}) = c + \tilde{\beta}_t \left(E[w'] - c + \int_0^{\bar{w}_{t+1}} F(w') dw' \right)$$

Clearly, $g_t(0) > 0$ and

$$\frac{\partial \bar{w}_t}{\partial \bar{w}_{t+1}} = g'_t(\bar{w}_{t+1}) = \tilde{\beta}_t F(\bar{w}_{t+1}) \in [0, 1).$$

The function $g_t(\bar{w}_{t+1})$ is therefore increasing in \bar{w}_{t+1} , and also convex because the slope is increasing. Also, since $\tilde{\beta}_t$ is decreasing in t , we have that $g_t(0)$ is decreasing in t and the slope $g'_t(\bar{w}_{t+1})$ is decreasing in t for every \bar{w}_{t+1} .

We can now proceed backwards, establishing that given an optimal \bar{w}_{t+1} , we show $\bar{w}_t > \bar{w}_{t+1}$. The proof is constructed using Figure 2.5. The reservation wage in the terminal period is equal to $\bar{w}_T = c$. Now set $t = T - 1$. The black solid line shows that the reservation wage \bar{w}_{T-1} in period $T - 1$ must satisfy $\bar{w}_T < \bar{w}_{T-1} < \tilde{w}_{T-1}$.

Now take $s = T - 2$. The red dashed line then implies that since $\bar{w}_{T-1} < \tilde{w}_{T-1}$, it must be that $\bar{w}_{T-1} < \bar{w}_{T-2} < \tilde{w}_{T-2}$.

Iterating this argument, we establish $\bar{w}_s > \bar{w}_t$ for $s < t$. Further, as $t \rightarrow -\infty$, we obtain $\bar{w}_t \rightarrow \bar{w}$. This completes the proof.

2.2.4 Backward induction

The preceding discussion illustrates an algorithm called the **backward induction**. The principle is to determine the decision rule in the terminal period T , which is easy, and then iterate backward on the recursion (2.16), obtaining the pairs of decision rules and values $(\bar{w}_t, V_t(w))$ for $t = T - 1, T - 2, \dots, 0$.

Observe that since the solution converges to the time-invariant problem when T is large, we can also use the backward induction algorithm to find the fixed point of the infinite-horizon problem (2.3).

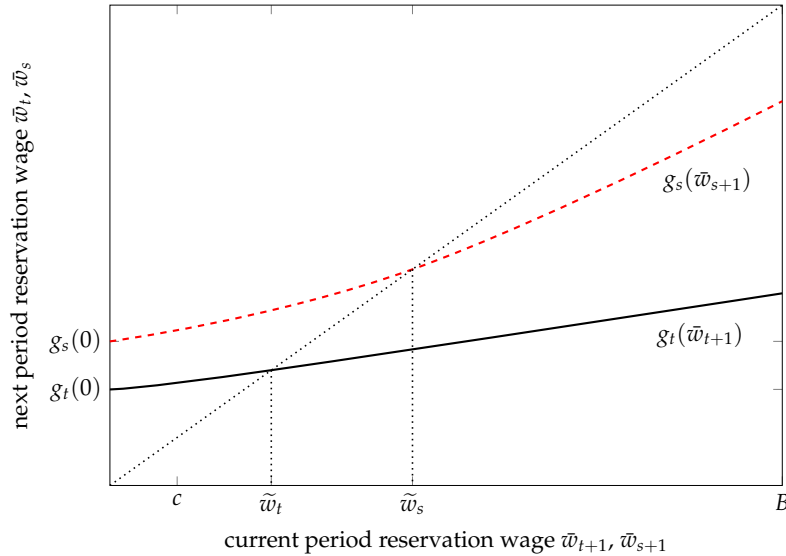


Figure 2.5: Optimal decision rules in the [McCall \(1970\)](#) model with a finite horizon. The graph plots decision rules for the current period reservation wage \bar{w}_t (\bar{w}_s) as a function of the next period reservation wage \bar{w}_{t+1} (\bar{w}_{s+1}) for two periods $s < t$.

Is solving the finite-horizon problem easier or harder than the infinite horizon problem? On the one hand, it is easier since the backward induction algorithm is straightforward and does not require calculating the fixed point. At the same time, we obtain a time-dependent decision rule which may become more memory intensive if all these rules need to be stored.

2.2.5 Welfare in a model with exogenous transitions

We now modify the setup, and work with the example of the lake model introduced in [Section 1.1.10](#). A worker in the economy can be either employed or unemployed. At the beginning of each period, a previously unemployed worker receives a wage offer with probability $\lambda > 0$ and accepts it, even if it is not optimal. The wage offer for an unemployed worker is drawn from a distribution with cdf $F(w)$ that has a density $f(w)$ with a full support on $[0, 1]$. Also at the beginning of each period, a previously employed worker separates from a job with probability $\delta > 0$.

Are all employed workers better off than the unemployed worker? To see whether this is the case, we need to compare the value functions for both types. In order to do so, we assume, in line with the previous setup, that workers are risk-neutral and discount future at rate $\beta \in (0, 1)$.

Denote V^u the value of an unemployed worker who has not received an offer at the beginning of the period, and $V^e(w)$ the value of an employed worker working at wage w . The Bellman equation for the unemployed worker is

$$V^u = c + \beta \left[(1 - \lambda) V^u + \lambda \int_0^1 V^e(w') dF(w') \right]$$

and for the employed worker,

$$V^e(w) = w + \beta [(1 - \delta) V^e(w) + \delta V^u].$$

The second equation can be used to express $V^e(w)$

$$V^e(w) = \frac{w + \beta \delta V^u}{1 - \beta(1 - \delta)}. \quad (2.18)$$

This can be substituted into the Bellman equation for V^u :

$$\begin{aligned} V^u &= c + \beta \left[(1 - \lambda) V^u + \lambda \int_0^1 \frac{w' + \beta \delta V^u}{1 - \beta(1 - \delta)} dF(w') \right] \\ &= c + \frac{\beta(1 - \lambda) + \beta^2 \lambda \delta - \beta^2(1 - \lambda)(1 - \delta)}{1 - \beta(1 - \delta)} V^u + \frac{\beta \lambda}{1 - \beta(1 - \delta)} E[w'] \end{aligned}$$

which can be solved for V^u :

$$V^u = \frac{(1 - \beta(1 - \delta))c + \beta \lambda E[w']}{1 + \beta(1 - \beta)(\lambda + \delta) + \beta^2}$$

Consequently, $V^e(w)$ can be expressed from (2.18) as

$$\begin{aligned} V^e(w) &= \frac{w}{1 - \beta(1 - \delta)} + \frac{\beta \delta c}{1 + \beta(1 - \beta)(\lambda + \delta) + \beta^2} \\ &\quad + \frac{\beta \delta}{1 - \beta(1 - \delta)} \frac{\beta \lambda E[w']}{1 + \beta(1 - \beta)(\lambda + \delta) + \beta^2}. \end{aligned}$$

We now need to compare V^u and $V^e(w)$. We are asking what are the levels of wages w such that $V^e(w) < V^u$. Substituting in, we obtain

$$w < \frac{(1 - \beta) [(1 - \beta(1 - \delta))c + \beta \lambda E[w']]}{1 + \beta(1 - \beta)(\lambda + \delta) + \beta^2}. \quad (2.19)$$

Observe that the right-hand side is strictly positive even in a situation when $c = 0$, when the unemployed worker receives no unemployment benefits. Hence some employed workers working at low wages are worse off than unemployed workers.

Why is a worker employed at a positive wage below the threshold given by (2.19) worse off than an unemployed worker with zero unemployment benefits? The unemployed worker receives no benefits but only waits to receive an offer from distribution $F(w)$. A currently employed worker with a very low w currently works at a very low wage until she is separated, after which she receives no unemployment benefits, and only then receives a typical wage offer from $F(w)$. This can be a worse path than starting from the unemployment state.

The reason is that separation is exogenous, the worker cannot choose to leave a job with a low wage, and she is stuck with it until the separation shock arrives. Workers with

these low wages exist under the stationary distribution because workers cannot choose to reject a low offer.

2.2.6 Further extensions

One can envision a whole variety of further extensions of this problem (see exercises in [Ljungqvist and Sargent \(2018\)](#), Chapter 6).

1. Allow the worker to **search on the job**. What does it imply for the reservation wage? How does the answer depend on whether wage offer distributions for the unemployed and currently employed differ? Could the worker potentially accept offers lower than c ?
2. Allow **seasonality** where the wage offer distribution may depend on the time period (for instance, over the business cycle). Is the value of the quit option still zero in such a case?
3. Assume that the worker does not know the wage offer distribution and must **learn** about it by sampling offers over time. See [Ljungqvist and Sargent \(2018\)](#), Section 6.6.
4. Worker can receive multiple offers per period, offers arrive probabilistically, or workers can choose how many costly offers to accept.

2.3 Equilibrium wage distribution

In this whole section, we assumed that the wage offer distribution is determined by nature. In reality, the offers are made by firms that have strategic considerations. We will later consider environments where workers receive offers and strategically accept or reject them, while firms compete for these workers and make these offers strategically in order to maximize profits.

One important and challenging question in this environment is why would firms want to make wage offers higher than the reservation wage \bar{w} when the worker would accept \bar{w} and any offer above \bar{w} reduces profits. This is the essence of the **Diamond paradox** ([Diamond \(1971\)](#)). Assume firms post wages with distribution $F(w)$ and the reservation wage is \bar{w} . But then firms have no incentives to post wages above \bar{w} and truncate $F(w)$ from above. This consequently reduces the reservation wage \bar{w} . This process continues until all firms post just the monopsony wage c which also becomes the reservation wage.

The key aspect of one way how to resolve this paradox is an element of imperfect information where the firm is uncertain how many or which workers it may attract, and offering a higher wage increases its chance to hire a worker. [Burdett and Judd \(1983\)](#) developed this idea in the context of a consumer goods market with non-sequential search where workers differ in terms of how many offers they receive in a given period. Since the firm does not know whether it faces a worker with less or more offers, it can aim at offering a higher wage, which increases chances of hiring but reduces profits conditional on hiring,

or, alternatively, at offering a low wage, reducing the chances of hiring but increasing profits conditional on hiring. The equilibrium outcome will then be a randomization strategy where some firms post high wages while others post low wages.

[Burdett and Mortensen \(1998\)](#)² offer a resolution in the context of the labor market with on-the-job search where firms can meet either unemployed or employed workers. These two groups differ in terms of their reservation wages. A similar randomization strategy then emerges. [Albrecht \(2011\)](#) provides a pedagogical treatment of this problem.

2.4 Numerical implementation

Solving the baseline [McCall \(1970\)](#) model is straightforward. Given parameters c, β, B and a distribution function $F(w)$, we can solve the algebraic equation (2.6) for \bar{w} :

$$\bar{w} - c = \frac{\beta}{1 - \beta} \int_{\bar{w}}^B (w' - \bar{w}) dF(w').$$

Given the analysis from Section 2.1.4, the difference between the left-hand side and right-hand side is monotone and crosses zero exactly once. The value of \bar{w} can therefore be found by interval bisection or a similar method.

2.4.1 Iteration on the reservation wage

Rather than using interval bisection, we can utilize the power of dynamic programming to find \bar{w} . As in equation (2.4), define

$$Q = c + \beta \int_0^B V(w') dF(w') = \frac{\bar{w}}{1 - \beta}$$

to be the value of obtaining an offer equal to the reservation wage. Then (2.3) can be written as

$$V(w) = \max_{\{\text{accept, reject}\}} \left\{ \frac{w}{1 - \beta}, Q \right\},$$

and plugging this expression into the previous equation yields

$$Q = c + \beta \int_0^B \max_{\{\text{accept, reject}\}} \left\{ \frac{w}{1 - \beta}, Q \right\} dF(w'). \quad (2.20)$$

This is again a type of a Bellman equation but now we are looking for a fixed point in the form of a scalar value $Q \in [0, (1 - \beta)^{-1} B]$. It can be shown that sufficient conditions for the application of the Banach fixed point theorem hold, and hence the following scheme will uncover a unique solution Q . Start with a guess $Q_0 \in [0, (1 - \beta)^{-1} B]$ and then iterate

²Despite the publication date, the ideas in this paper existed already in the 1980s.

on

$$Q_{n+1} = c + \beta \int_0^B \max_{\{\text{accept}, \text{reject}\}} \left\{ \frac{w}{1-\beta}, Q_n \right\} dF(w').$$

The fixed point theorem implies that $Q_n \rightarrow Q$, and hence the reservation wage can be obtained as $\bar{w} = \lim_{n \rightarrow \infty} (1-\beta) Q_n$.

2.4.2 Value function iteration

However, there is also a more general approach that utilizes the method of **successive approximation** or **value function iteration**. Consider again equation (2.3) and write it as follows:

$$V_{n+1}(w) = \max_{\{\text{accept}, \text{reject}\}} \left\{ \frac{w}{1-\beta}, c + \beta \int_0^B V_n(w') dF(w') \right\} \quad (2.21)$$

That is, given a function $V_n(w)$, it is conceptually straightforward to compute the function $V_{n+1}(w)$. We can start with an initial guess, say $V_0(w) = 0$, and see whether the algorithm converges numerically: $V_n(w) \rightarrow V_\infty(w) = V(w)$ as $n \rightarrow \infty$. In the subsequent chapters, we will formally establish conditions under which this approach is valid. The reservation wage \bar{w} is then determined as

$$\frac{\bar{w}}{1-\beta} = c + \beta \int_0^B V_\infty(w') dF(w').$$

Notice that successive approximations in (2.21) mimic the idea of **backward induction**. With $V_0(w) = 0$, $V_n(w)$ corresponds to the value of the finite-horizon problem with n periods remaining. The limit $n \rightarrow \infty$ corresponds to the idea that the value in a finite-horizon problem with an increasing horizon converges to the fixed point in the infinite-horizon problem.

The second aspect of the numerical solution concerns encoding the value function on a computer. Equation (2.21) involves a function of the wage w , and an integral over $F(w)$. When the wage distribution $F(w)$ only has a finite number of mass points with probabilities f^i at wages w^i , $i = 0, \dots, I$ such that $w^0 < w^1 < \dots < w^I$, then we only need to store $I+1$ values $V^i = V(w^i)$, and integration is replaced by a sum. We can then replace the functional equation (2.21) with the algebraic system

$$V_{n+1}^i = \max_{\{\text{accept}, \text{reject}\}} \left\{ \frac{w^i}{1-\beta}, c + \beta \sum_{j=0}^I V_n^j f^j \right\} \quad i = 0, \dots, I$$

We are then looking for a vector $V = (V^0, \dots, V^I)$ that is the fixed point of this discrete problem, i.e., a V that solves

$$V^i = \max_{\{\text{accept}, \text{reject}\}} \left\{ \frac{w^i}{1-\beta}, c + \beta \sum_{j=0}^I V^j f^j \right\} \quad i = 0, \dots, I.$$

When $\beta \in (0, 1)$, the conditions of the Banach fixed point theorem hold for this problem.

Hence a unique fixed point V exists and $\lim_{n \rightarrow \infty} V_n = V$ regardless of the choice of the initial guess V^0 .

Notice that in the case of this discrete wage distribution, the assumption of a full support on $[0, B]$ that we made in Section 2.1.4 does not hold. The solution to equation (2.6) will typically lie in the interior of an interval (w^i, w^{i+1}) for some i , and then any wage in this interval can be interpreted as the reservation wage. However, this nonuniqueness does not have substantial consequences because the optimal policy (accept, reject, or be indifferent) for any wage offer that can be drawn with a strictly positive probability remains unique.

On the other hand, how should we proceed if the distribution of wage offers is continuous? One numerical approach is to discretize the problem on a grid, and replace the function $V(w)$ with a vector \hat{V} . Let us choose a grid of nodes w^i , $i = 0, \dots, I$, such that $0 = w^0 < w^1 < \dots < w^I = B$. One frequent choice is to split the grid equidistantly in I subintervals of length B/I , such that $w^i = Bi/I$. We then replace the distribution of wage offers $F(w)$ on $[0, B]$ with a discrete distribution $\hat{f}^i \doteq \hat{f}(w^i)$ on nodes w^i that approximates $F(w)$, for example,

$$\hat{f}^i = \begin{cases} F\left(\frac{1}{2}(w^1 + w^0)\right) & i = 0 \\ F\left(\frac{1}{2}(w^{i+1} + w^i)\right) - F\left(\frac{1}{2}(w^i + w^{i-1})\right) & 0 < i < I \\ 1 - F\left(\frac{1}{2}(w^I + w^{I-1})\right) & i = I \end{cases}$$

We now replaced the continuous wage offer distribution with a discretized counterpart and can further proceed as in the case of a discrete distribution. We replace the functional equation (2.21) with the algebraic system

$$\hat{V}_{n+1}^i = \max_{\{\text{accept, reject}\}} \left\{ \frac{w^i}{1 - \beta}, c + \beta \sum_{j=0}^I \hat{V}_n^j \hat{f}^j \right\} \quad i = 0, \dots, I$$

where $\hat{V}_n = (\hat{V}_n^0, \dots, \hat{V}_n^I)'$ is a vector with $I + 1$ elements. Given that the integral in (2.21) is approximated as

$$\int_0^B V_n(w') dF(w') \approx \sum_{i=0}^I V_n(w^i) \hat{f}^i,$$

we can also expect that $\hat{V}_n^i \approx V_n(w^i)$. As before, we are then looking for a vector \hat{V} that is the fixed point of the discrete problem, i.e., a \hat{V} that solves

$$\hat{V}^i = \max_{\{\text{accept, reject}\}} \left\{ \frac{w^i}{1 - \beta}, c + \beta \sum_{j=0}^I \hat{V}^j \hat{f}^j \right\} \quad i = 0, \dots, I.$$

Once the fixed point is found, we associate the values $V(w^i)$ on the grid with \hat{V}^i . The values $V(w)$ outside the gridpoints w^i then can be approximated using a desired interpo-

lation method.

There are two important considerations underlying this discretization. First, the discretized problem may have a unique solution even though the original problem may have zero or more than one solutions. This issue is primarily caused by the fact that fixed point results that establish existence of unique fixed points crucially rely on compactness properties of the underlying state space and the space of functions on which we look for solutions. Discretizing functions from a continuous state space to finite grid can change these compactness properties in profound ways. More details can be found in [Borovička and Stachurski \(2021\)](#).

Second, even if both the original and discretized problem have a unique solution, we would like to establish how ‘close’ the function $V(w)$ and the vector \hat{V} are. We do not establish a formal approximation result here but a large area in numerical mathematics studies these problems. [Judd \(1998\)](#) is a good starting point for a treatment aimed at economists.

2.4.3 Evaluating the expectations operator

The numerical implementation in Section 2.4.2 involves the approximation of the integral

$$\int_0^B V(w') dF(w') \approx \sum_{i=0}^I V(w^i) \hat{f}^i. \quad (2.22)$$

Such a discrete-state approximation is called a **quadrature rule**. In this approximation, we utilized the same grid $\{w^i\}_{i=0}^I$ that we used to construct the approximation of the value function. The method therefore only requires the initial construction of the probability mass distribution $\{\hat{f}^i\}_{i=0}^I$, and does not involve any interpolation during the iteration steps.

Interpolation

On the other hand, there may good reasons why the evaluation of the integral may rely on a different grid than the one on which we approximated the value function. For example, it may be desirable to keep the grid for \hat{V} sparse, while at the same time constructing a more accurate approximation of the integral on a finer or differently spaced grid. Gaussian quadrature, described in Section 2.4.4, is one such example.

Specifically, we are now designing a set of nodes $\{\tilde{w}^j\}_{j=1}^J$ and associated weights $\{\tilde{f}^j\}_{j=1}^J$ to evaluate

$$\int_0^B V(w') dF(w') \approx \sum_{j=1}^J V(\tilde{w}^j) \tilde{f}^j. \quad (2.23)$$

When the function $V(w)$ is known, then we can proceed as in Section 2.4.2. However, when the function is approximated using the vector of values \hat{V} on the set of nodes $\{w^i\}_{i=0}^I$,

then an additional step is required. Since the grid $\{\tilde{w}^j\}_{j=1}^J$ will typically not align with the the grid $\{w^i\}_{i=0}^I$ on which the discretization of the value function is constructed, the values $V(\tilde{w}^j)$ must be interpolated using a particular interpolation method using the known values $\hat{V}^i = \hat{V}(w^i)$. In the case of the value function iteration from Section 2.4.2, since V changes with every iteration, this interpolation has to be conducted in each iteration step.

The simplest form of interpolation is linear interpolation, which assumes that the function $V(w)$ is piecewise linear. Imagine node \tilde{w}^j lies in the interval $[w^i, w^{i+1}]$ for a particular value of $i = 0, \dots, I - 1$. Then $V(\tilde{w}^j)$ can be approximated as

$$V(\tilde{w}^j) = \hat{V}(w^i) + (\tilde{w}^j - w^i) \frac{\hat{V}(w^{i+1}) - \hat{V}(w^i)}{w^{i+1} - w^i}.$$

Software packages provide powerful interpolation methods that extend beyond linear interpolation.

Monte Carlo approach

An alternative to quadrature is to approximate the expectations operator using a simple Monte Carlo simulation. This approach is based on the Central Limit Theorem, which states that the empirical distribution of a large sample of independent draws from a particular random variable converges to the theoretical distribution.

Therefore, draw a large number of draws $\{\tilde{w}^d\}_{d=1}^D$ from $F(w)$ and approximate

$$\int_0^B V(w') dF(w') = E[w] \approx \frac{1}{D} \sum_{d=1}^D V(\tilde{w}^d).$$

Notice the absence of the density function in the sum on the right-hand side. Here, the random number generator already accounts for the proper weighting of the draws, since values \tilde{w}^d in the parts of the state space that have a higher density will be drawn more frequently. Again, since \tilde{w}^d will typically not lie on the grid, values $V(\tilde{w}^d)$ need to be interpolated using a suitable algorithm.

2.4.4 Gaussian quadrature

Gaussian quadrature designs the choice of nodes $\{\tilde{w}^j\}_{j=1}^J$ and associated weights $\{\tilde{f}^j\}_{j=1}^J$ in the rule (2.23) in a particular way to yield a good approximation for a particular class of functions. Specifically, a J node approximation is constructed to provide an exact formula for the evaluation of the expectation of all polynomial functions up to degree $2J - 1$. This also means that such quadrature will provide a good approximation to function $V(w)$ if the function can be well approximated using such polynomials. Naturally, the choice of nodes and weights depends on the particular distribution $F(w)$ as well on the integration range.

We first introduce general theory and then provide specific formulas for concrete cases.

Additional information can be found in Gil et al. (2007), Section 5.3, or in the [Wikipedia article on Gaussian quadrature](#).

We generalize problem (2.23) and consider the approximation of the integral of functions $V(w)$

$$\int_{\underline{w}}^{\bar{w}} V(w) f(w) dw \approx \sum_{j=1}^J V(\bar{w}^j) \tilde{f}^j$$

on interval (\underline{w}, \bar{w}) under a weighting function $f(w)$ which may not be a density. The idea is to construct an orthogonal basis of polynomials $P_n(w)$, $n = 0, 1, \dots, J$ of the space of all polynomials of degree up to J . The polynomials in the basis are orthogonal to each other under the inner product implied by the weighting function $f(w)$:

$$\langle P_m, P_n \rangle \doteq \int_{\underline{w}}^{\bar{w}} P_m(w) P_n(w) f(w) dw = 0 \quad m \neq n. \quad (2.24)$$

We normalize the zero-th degree polynomial to $P_0(w) = 1$. The orthogonality restriction (2.24) for $m = 0$ also means that

$$\langle 1, P_n \rangle = \int_{\underline{w}}^{\bar{w}} P_n(w) f(w) dw = 0 \quad n \geq 1.$$

Since the polynomials $P_n(w)$ form a basis of the space of polynomials of degree up to J , any polynomial $h(w)$ of degree up to J can be written as a linear combination of $P_n(w)$, $n = 0, \dots, J$.

The idea of the J -node Gaussian quadrature approximation is to choose the nodes \bar{w}^j , $j = 1, \dots, J$ to be the zeros of polynomial $P_J(w)$. The following theorem formalizes the sense in which this is a good choice. The resulting set of nodes and associated weights is such that the quadrature rule exactly evaluates the integral of any polynomial of degree up to $2J - 1$, so it will approximate a given function well as long as such function can be well approximated using such a polynomial.

Theorem 2.1. *Let $P_n(w)$, $n = 0, 1, \dots, J$ be the orthogonal basis of the space of polynomials of degree up to J on $[\underline{w}, \bar{w}]$ under a weighting function $f(w)$, and \bar{w}^j , $j = 1, \dots, J$ the roots of $P_J(w)$. Then there exist weights \tilde{f}^j , $j = 1, \dots, J$ such that the quadrature rule*

$$\int_{\underline{w}}^{\bar{w}} h(w) f(w) dw = \sum_{j=1}^J h(\bar{w}^j) \tilde{f}^j$$

is exact for all polynomials of degree up to $2J - 1$. Moreover, all the nodes \bar{w}^j lie in the open interval (\underline{w}, \bar{w}) .

Proof. Let $h(w)$ be a polynomial of degree up to $2J - 1$. Dividing this polynomial by $P_J(w)$ yields the decomposition

$$h(w) = q(w) P_J(w) + r(w)$$

where the quotient $q(w)$ and the remainder $r(w)$ are polynomials of degree at most $J - 1$. This also means that $q(w)$ and $r(w)$ can be written as linear combinations of $P_n(w)$, $n = 0, \dots, J - 1$, and are therefore orthogonal to $P_J(w)$. Therefore also

$$\begin{aligned} \int_{\underline{w}}^{\bar{w}} h(w) f(w) dw &= \int_{\underline{w}}^{\bar{w}} q(w) P_J(w) f(w) dw + \int_{\underline{w}}^{\bar{w}} r(w) f(w) dw \\ &= \int_{\underline{w}}^{\bar{w}} r(w) f(w) dw. \end{aligned} \quad (2.25)$$

Now take the nodes $\tilde{w}^j, j = 1, \dots, J$ and construct the basis of Lagrange polynomials

$$l_j(w) = \prod_{\substack{i=1 \\ i \neq j}}^J \frac{w - \tilde{w}^i}{\tilde{w}^j - \tilde{w}^i}. \quad (2.26)$$

Notice that the Lagrange polynomial $l_j(w)$ satisfies $l_j(\tilde{w}^j) = 1$ and $l_j(\tilde{w}^i) = 0$ for $i \neq j$. Then we can write the function $r(w)$ as

$$r(w) = \sum_{j=1}^J r(\tilde{w}^j) l_j(w).$$

To see, this, notice that the left-hand and right-hand sides certainly coincide at the nodes $\tilde{w}^j, j = 1, \dots, J$. Since both sides are polynomials of degree at most $J - 1$ that coincide at J points, they must be identical everywhere. The right-hand side is the Lagrange interpolating polynomial, the unique polynomial of the lowest degree that interpolates the given data $(\tilde{w}^j, r(\tilde{w}^j)), j = 1, \dots, J$.

We can then write

$$\int_{\underline{w}}^{\bar{w}} r(w) f(w) dw = \sum_{j=1}^J r(\tilde{w}^j) \underbrace{\int_{\underline{w}}^{\bar{w}} l_j(w) f(w) dw}_{= \tilde{f}^j} \quad (2.27)$$

where the weights \tilde{f}^j are defined as the integrals of Lagrange polynomials, and they are independent of the choice of the function $h(w)$.

Further, notice that the functions $h(w)$ and $r(w)$ coincide at the nodes \tilde{w}^j because $P_J(\tilde{w}^j) = 0$:

$$h(\tilde{w}^j) = q(\tilde{w}^j) P_J(\tilde{w}^j) + r(\tilde{w}^j) = r(\tilde{w}^j).$$

Consequently, we can combine (2.25) and (2.27), and write

$$\int_{\underline{w}}^{\bar{w}} h(w) f(w) dw = \int_{\underline{w}}^{\bar{w}} r(w) f(w) dw = \sum_{j=1}^J r(\tilde{w}^j) \tilde{f}^j = \sum_{j=1}^J h(\tilde{w}^j) \tilde{f}^j.$$

So indeed, the quadrature formula

$$\int_{\underline{w}}^{\overline{w}} h(w) f(w) dw = \sum_{j=1}^J h(\tilde{w}^j) \tilde{f}^j$$

with nodes \tilde{w}^j given by the zeros of the polynomial $P_J(w)$ and weights given by

$$\tilde{f}^j = \int_{\underline{w}}^{\overline{w}} l_j(w) f(w) dw \quad (2.28)$$

is indeed exact for all polynomials of degree up to $2J - 1$.

It remains to be shown that the nodes \tilde{w}^j corresponding to the zeros of $P_J(w)$ lie in the open interval $(\underline{w}, \overline{w})$. Let us consider the consequences that alternative roots have on the sign of $P_J(w)$ on $(\underline{w}, \overline{w})$ by decomposing $P_J(w)$ into the product of its factors. The factors associated with complex roots, roots outside the interval $(\underline{w}, \overline{w})$, and roots with even multiplicity do not change the sign of $P_J(w)$ on $(\underline{w}, \overline{w})$. What remains are roots on $(\underline{w}, \overline{w})$ with odd multiplicity. Denote those roots w^i and multiply $P_J(w)$ by $(w - w^i)$ for each occurrence of such a root to form

$$P_J(w) \prod_i (w - w^i),$$

where the product is over all the roots w^i with odd multiplicity on $(\underline{w}, \overline{w})$. This resulting polynomial does not change sign on $(\underline{w}, \overline{w})$, so that

$$\int_{\underline{w}}^{\overline{w}} P_J(w) \prod_i (w - w^i) f(w) dw \neq 0$$

because the weighting function is nonnegative as well. But since $P_J(w)$ is orthogonal to all polynomials of order up to $J - 1$, it must be that $\prod_i (w - w^i)$ is a polynomial of order at least J , hence $P_J(w)$ has exactly J distinct roots on the interval $(\underline{w}, \overline{w})$. ■

There are many alternative methods how to find the roots of the polynomial $P_J(w)$, some of them extremely efficient even for very large values of J . We discuss one easily implementable method using specific cases of Gaussian quadrature associated with particular weighting functions and integration intervals $(\underline{w}, \overline{w})$.

Quadrature weights

Before we do so, we discuss in more detail the construction of the associated weights \tilde{f}^j in (2.28). Write the polynomials $P_n(w)$ as

$$P_n(x) = a_{n,0} + a_{n,1}x + \dots + a_{n,n}x^n.$$

Consider the numerators of Lagrange polynomials $l_j(w)$ in (2.26)

$$\prod_{\substack{i=1 \\ i \neq j}}^J (w - \tilde{w}^i) = \frac{\prod_{i=1}^J (w - \tilde{w}^i)}{w - \tilde{w}^j} = \frac{P_J(w)}{a_{J,J}(w - \tilde{w}^j)}.$$

Taking the limit as $w \rightarrow \tilde{w}^j$, applying the L'Hospital rule to the right-hand side, yields

$$\prod_{\substack{i=1 \\ i \neq j}}^J (\tilde{w}^j - \tilde{w}^i) = \frac{P_J'(\tilde{w}^j)}{a_{J,J}},$$

which is the denominator of the Lagrange polynomial $l_j(w)$ in (2.26). We can thus combine the last two expressions to write the weight (2.28) as

$$\tilde{f}^j = \int_{\underline{w}}^{\bar{w}} l_j(w) f(w) dw = \frac{1}{P_J'(\tilde{w}^j)} \int_{\underline{w}}^{\bar{w}} \frac{P_J(w)}{w - \tilde{w}^j} f(w) dw. \quad (2.29)$$

Since

$$\frac{(\tilde{w}^j)^k}{w - \tilde{w}^j} = \frac{(\tilde{w}^j)^k - w^k}{w - \tilde{w}^j} + \frac{w^k}{w - \tilde{w}^j},$$

we have, for $k \leq J$,

$$\begin{aligned} \int_{\underline{w}}^{\bar{w}} \frac{P_J(w) w^k}{w - \tilde{w}^j} f(w) dw &= \int_{\underline{w}}^{\bar{w}} \frac{P_J(w) (\tilde{w}^j)^k}{w - \tilde{w}^j} f(w) dw - \int_{\underline{w}}^{\bar{w}} P_J(w) \frac{(\tilde{w}^j)^k - w^k}{w - \tilde{w}^j} f(w) dw \\ &= (\tilde{w}^j)^k \int_{\underline{w}}^{\bar{w}} \frac{P_J(w)}{w - \tilde{w}^j} f(w) dw \end{aligned}$$

because the last expression on the first line is the inner product of $P_J(w)$ with a polynomial of degree less than J , which is orthogonal to $P_J(w)$ by construction, and hence the inner product is zero. Since this expression holds for any w^k , $k \leq J$, it also holds for any polynomial $q(w)$ of degree at most J :

$$\int_{\underline{w}}^{\bar{w}} \frac{P_J(w) q(w)}{w - \tilde{w}^j} f(w) dw = q(\tilde{w}^j) \int_{\underline{w}}^{\bar{w}} \frac{P_J(w)}{w - \tilde{w}^j} f(w) dw.$$

Notice that we can write

$$\frac{P_J(w)}{w - \tilde{w}^j} = a_{J,J} w^{J-1} + s(w),$$

where $s(w)$ is a polynomial of degree at most $J - 2$. Now choose $q(w) = P_{J-1}(w)$. Then

$$\begin{aligned} \int_{\underline{w}}^{\bar{w}} \frac{P_J(w)}{w - \tilde{w}^j} f(w) dw &= \frac{1}{P_{J-1}(\tilde{w}^j)} \int_{\underline{w}}^{\bar{w}} \frac{P_J(w) P_{J-1}(w)}{w - \tilde{w}^j} f(w) dw \\ &= \frac{a_{J,J}}{P_{J-1}(\tilde{w}^j)} \int_{\underline{w}}^{\bar{w}} w^{J-1} P_{J-1}(w) f(w) dw \end{aligned}$$

where we used the decomposition of $P_J(w) / (w - \tilde{w}^j)$, noticing that $s(w)$ is orthogonal to $P_{J-1}(w)$. Using a similar argument, write

$$w^{J-1} = \left(w^{J-1} - \frac{P_{J-1}(w)}{a_{J-1,J-1}} \right) + \frac{P_{J-1}(w)}{a_{J-1,J-1}}$$

where the term in parenthesis is a polynomial of degree at most $J - 2$, and hence orthogonal to $P_{J-1}(w)$, which allows us to write the previous expression as

$$\begin{aligned} \int_{\underline{w}}^{\bar{w}} \frac{P_J(w)}{w - \tilde{w}^j} f(w) dw &= \frac{a_{J,J}}{P_{J-1}(\tilde{w}^j)} \int_{\underline{w}}^{\bar{w}} \frac{P_{J-1}(w)}{a_{J-1,J-1}} P_{J-1}(w) f(w) dw \\ &= \frac{a_{J,J}}{a_{J-1,J-1} P_{J-1}(\tilde{w}^j)} \int_{\underline{w}}^{\bar{w}} P_{J-1}(w)^2 f(w) dw \end{aligned}$$

The weight \tilde{f}^i in (2.29) can therefore be written as

$$\tilde{f}^j = \int_{\underline{w}}^{\bar{w}} l_j(w) f(w) dw = \frac{a_{J,J}}{a_{J-1,J-1}} \frac{\int_{\underline{w}}^{\bar{w}} P_{J-1}(w)^2 f(w) dw}{P'_J(\tilde{w}^j) P_{J-1}(\tilde{w}^j)}$$

The weights \tilde{f}^j are strictly positive. To see this, construct the polynomial

$$g_j(w) = l_j(w)^2 = \prod_{\substack{i=1 \\ i \neq j}}^J \frac{(w - \tilde{w}^i)^2}{(\tilde{w}^j - \tilde{w}^i)^2}.$$

Notice that $g_j(\tilde{w}^i) = \delta_{ij} = 1_{\{i=j\}}$. Since $g_j(w)$ is a polynomial of degree less than $2J - 1$, the quadrature formula evaluates its integral under weight $f(w)$ exactly

$$0 < \int_{\underline{w}}^{\bar{w}} g_i(w) f(w) dw = \sum_{j=1}^J g_i(\tilde{w}^j) \tilde{f}^j = \sum_{j=1}^J \delta_{ij} \tilde{f}^j = \tilde{f}^i, \quad i = 1, \dots, J.$$

Construction of orthogonal polynomial base

The orthogonal polynomial base defined using the scalar product $\langle \cdot, \cdot \rangle$ in (2.24) with polynomials $P_n(w)$ normalized to be monic, i.e., with leading coefficient $a_{n,n} = 1$, satisfies the recurrence relation

$$P_{r+1}(w) = (w - a_{r,r}) P_r(w) - a_{r,r-1} P_{r-1}(w) - \dots - a_{r,0} P_0(w). \quad (2.30)$$

The coefficients are given by

$$a_{r,s} = \frac{\langle w P_r, P_s \rangle}{\langle P_s, P_s \rangle}. \quad (2.31)$$

This relation, known as the Gram–Schmidt orthogonalization process, can be proven by induction. Start with $P_0(w) = 1$. For $r = 0$, expression (2.30) yields

$$P_1(w) = (w - a_{0,0}) P_0(w)$$

where

$$a_{0,0} = \frac{\langle wP_0, P_0 \rangle}{\langle P_0, P_0 \rangle} = \frac{\int_{\underline{w}}^{\bar{w}} wf(w) dw}{\int_{\underline{w}}^{\bar{w}} f(w) dw}$$

Then we have

$$\begin{aligned} \langle P_1, P_0 \rangle &= \langle (w - a_{0,0}) P_0, P_0 \rangle = \langle wP_0, P_0 \rangle - a_{0,0} \langle P_0, P_0 \rangle \\ &= a_{0,0} \langle P_0, P_0 \rangle - a_{0,0} \langle P_0, P_0 \rangle = 0. \end{aligned}$$

Now, for $r \geq 1$, assume that polynomials $P_0(w), \dots, P_r(w)$ are orthogonal. Then for $s \leq r$,

$$\begin{aligned} \langle P_{r+1}, P_s \rangle &= \langle wP_r, P_s \rangle - a_{r,r} \langle P_r, P_s \rangle - a_{r,r-1} \langle P_{r-1}, P_s \rangle - \dots - a_{r,0} \langle P_0, P_s \rangle \\ &= \langle wP_r, P_s \rangle - a_{r,s} \langle P_s, P_s \rangle \\ &= \langle wP_r, P_s \rangle - \frac{\langle wP_r, P_s \rangle}{\langle P_s, P_s \rangle} \langle P_s, P_s \rangle = 0, \end{aligned}$$

where the second equality follows from the fact that $\langle P_m, P_s \rangle = 0$ for $m \neq s$.

The recurrence relation can be reduced to just three adjacent terms because the scalar product satisfies $\langle wf, g \rangle$ for any two functions $f(w), g(w)$. For $s < r - 1$, the polynomial $wP_s(w)$ is of order less than r , and hence orthogonal to P_r . This implies

$$a_{r,s} = \frac{\langle wP_r, P_s \rangle}{\langle P_s, P_s \rangle} = \frac{\langle P_r, wP_s \rangle}{\langle P_s, P_s \rangle} = 0.$$

Then the recurrence relation (2.30) reduces to

$$P_{r+1}(w) = (w - a_{r,r}) P_r(w) - a_{r,r-1} P_{r-1}(w) \quad (2.32)$$

with the convention $P_{-1}(w) = 0$. Finally, notice that

$$a_{r,r-1} = \frac{\langle wP_r, P_{r-1} \rangle}{\langle P_{r-1}, P_{r-1} \rangle} = \frac{\langle P_r, wP_{r-1} \rangle}{\langle P_{r-1}, P_{r-1} \rangle} = \frac{\langle P_r, P_r \rangle}{\langle P_{r-1}, P_{r-1} \rangle}$$

because

$$\langle P_r, wP_{r-1} \rangle = \langle P_r, P_r \rangle + \langle P_r, wP_{r-1} - P_r \rangle = \langle P_r, P_r \rangle$$

because the polynomial $wP_{r-1} - P_r$ is of order at most $r - 1$, and hence orthogonal to P_r .

Normalization of the polynomial base

In what follows, it will be useful to work with an orthonormal polynomial base in which all base polynomials are normalized to have unit norm. This normalization can be con-

structured in a simple way by defining $\lambda_r = \sqrt{\langle P_r, P_r \rangle}$, and constructing normalized polynomials

$$\hat{P}_r(w) = \lambda_r^{-1} P_r(w)$$

which satisfy $\|\hat{P}_r\| = \sqrt{\langle \hat{P}_r, \hat{P}_r \rangle} = 1$. Then we can rewrite (2.32) as

$$\lambda_{r+1} \hat{P}_{r+1}(w) = (w - a_{r,r}) \lambda_r \hat{P}_r(w) - a_{r,r-1} \lambda_{r-1} \hat{P}_{r-1}(w).$$

Using the definitions of $a_{r,r}$ and $a_{r,r-1}$,

$$a_{r,r} = \frac{\langle w P_r, P_r \rangle}{\langle P_r, P_r \rangle} \quad a_{r,r-1} = \frac{\langle P_r, P_r \rangle}{\langle P_{r-1}, P_{r-1} \rangle}$$

allows us to rewrite the recurrence relation as

$$\sqrt{a_{r+1,r}} \hat{P}_{r+1}(w) = (w - a_{r,r}) \hat{P}_r(w) - \sqrt{a_{r,r-1}} \hat{P}_{r-1}(w). \quad (2.33)$$

Given that $P_r(w)$ is monic, the leading coefficient in $\hat{P}_r(w)$ is equal to $\hat{a}_{r,r} = \langle P_r, P_r \rangle^{-1/2}$. The coefficients can be rewritten in terms of the normalized polynomials as

$$a_{r,r} = \langle w \hat{P}_r, \hat{P}_r \rangle \quad \sqrt{a_{r,r-1}} = \frac{\hat{a}_{r-1,r-1}}{\hat{a}_{r,r}}$$

Initial terms in the series of polynomials are defined as $\hat{P}_{-1}(w) = 0$ and

$$\hat{P}_0(w) = \frac{P_0(w)}{\sqrt{\langle P_0, P_0 \rangle}} = \left(\int_{\underline{w}}^{\bar{w}} f(w) dw \right)^{-1/2} = \frac{1}{\sqrt{\mu_0}}.$$

The coefficients of the normed polynomials can be obtained by renormalizing the monic coefficients obtained from recursion (2.32).

Golub–Welsch algorithm

The Golub–Welsch algorithm transforms the problem of finding the roots of the polynomial $P_J(w)$ into a problem of finding the eigenvalues of a particular tridiagonal matrix. Denote

$$\pi_{J-1}(w) = (P_0(w), \dots, P_{J-1}(w))'$$

the column vector of polynomials, e_j the coordinate vector $e_j = (0, \dots, 0, 1)'$, and

$$\Lambda = \begin{pmatrix} a_{0,0} & 1 & 0 & \dots & \dots & \dots \\ a_{1,0} & a_{1,1} & 1 & 0 & \dots & \dots \\ 0 & a_{2,1} & a_{2,2} & 1 & 0 & \dots \\ 0 & \dots & \dots & \dots & \dots & 0 \\ \dots & \dots & 0 & a_{J-2,J-3} & a_{J-2,J-2} & 1 \\ \dots & \dots & \dots & 0 & a_{J-1,J-2} & a_{J-1,J-1} \end{pmatrix}$$

the tridiagonal matrix called the Jacobi matrix. Then the three-term recurrence equation (2.32) can be written in the form

$$\Lambda \pi_{J-1}(w) = w \pi_{J-1}(w) - P_J(w) e_J. \quad (2.34)$$

This implies that the roots of the polynomial $P_J(w)$ are identical to the eigenvalues of the Jacobi matrix Λ . To see this, observe that the column vector $P_J(w) e_J$ is equal to zero if and only if w is one of the roots of $P_J(w)$. For these roots \tilde{w}^j , $j = 1, \dots, J$, the recurrence equation (2.34) reduces to

$$\Lambda \pi_{J-1}(\tilde{w}^j) = \tilde{w}^j \pi_{J-1}(\tilde{w}^j),$$

which is the eigenvector equation for the matrix Λ , with eigenvalues \tilde{w}^j and eigenvectors $\pi_{J-1}(\tilde{w}^j)$, $j = 1, \dots, J$. Recall that the vector $\pi_{J-1}(\tilde{w}^j)$ is the vector of the values of polynomials $P_n(w)$, $n = 0, \dots, J-1$ evaluated at the value \tilde{w}^j , which is a root of polynomial $P_J(w)$.

Instead of solving the eigenvalue problem for Λ , it is advantageous to solve for the eigenvalues of the symmetric matrix $\hat{\Lambda}$ defined using the recurrence relation on the orthonormal polynomial base (2.33). Defining

$$\hat{\pi}_{J-1}(w) = (\hat{P}_0(w), \dots, \hat{P}_{J-1}(w))'$$

and

$$\hat{\Lambda} = \begin{pmatrix} a_{0,0} & \sqrt{a_{1,0}} & 0 & \dots & \dots & \dots \\ \sqrt{a_{1,0}} & a_{1,1} & \sqrt{a_{2,1}} & 0 & \dots & \dots \\ 0 & \sqrt{a_{2,1}} & a_{2,2} & \sqrt{a_{3,2}} & 0 & \dots \\ 0 & \dots & \dots & \dots & \dots & 0 \\ \dots & \dots & 0 & \sqrt{a_{J-2,J-3}} & a_{J-2,J-2} & \sqrt{a_{J-1,J-2}} \\ \dots & \dots & \dots & 0 & \sqrt{a_{J-1,J-2}} & a_{J-1,J-1} \end{pmatrix} \quad (2.35)$$

the recurrence relation for the orthonormal polynomial base can be written as

$$\hat{\Lambda} \hat{\pi}_{J-1}(w) = w \hat{\pi}_{J-1}(w) - \sqrt{a_{J,J-1}} \hat{P}_J(w) e_J.$$

The polynomial $\hat{P}_J(w)$ has the same roots \tilde{w}^j as the monic polynomial $P_J(w)$, and for these roots, we have

$$\hat{\Lambda} \hat{\pi}_{J-1}(\tilde{w}^j) = \tilde{w}^j \hat{\pi}_{J-1}(\tilde{w}^j),$$

so \tilde{w}^j and $\hat{\pi}_{J-1}(\tilde{w}^j)$ are eigenvalues and eigenvectors of the Jacobi matrix $\hat{\Lambda}$. The matrices Λ and $\hat{\Lambda}$ have identical eigenvalues but the eigenvectors differ. Identity of the eigenvalues also comes from the fact that Λ and $\hat{\Lambda}$ are similar matrices, meaning that there exists an invertible matrix P such that $\hat{\Lambda} = P^{-1} \Lambda P$.

Finally, we determine the weights \tilde{f}^j . Since $\hat{P}_i \hat{P}_k$, $i, k = 0, \dots, n-1$ are polynomials of degree smaller than $2J-1$, the integrals

$$\langle \hat{P}_i, \hat{P}_k \rangle = \int_w^{\bar{w}} \hat{P}_i(w) \hat{P}_k(w) f(w) dw$$

can be exactly expressed using the quadrature formula

$$\delta_{ik} = \langle \hat{P}_i, \hat{P}_k \rangle = \sum_{j=1}^J \hat{P}_i(\tilde{w}^j) \hat{P}_k(\tilde{w}^j) \tilde{f}^j.$$

This set of equations can be expressed in matrix form as

$$\Pi' W \Pi = I$$

where $W = \text{diag}(\tilde{f}^1, \dots, \tilde{f}^J)$ is the diagonal matrix of quadrature weights, and $\Pi = (\hat{\pi}_{J-1}(\tilde{w}^1), \dots, \hat{\pi}_{J-1}(\tilde{w}^J))'$ is the matrix the polynomial vectors $\hat{\pi}_{J-1}(\tilde{w}^i)$ in its rows. These polynomial vectors are also eigenvectors of the matrix $\hat{\Lambda}$. Given that W and I are full rank, the matrix Π is invertible, and

$$W^{-1} = \Pi \Pi'.$$

The i -th element of the diagonal is

$$\frac{1}{\tilde{f}^i} = \hat{\pi}_{J-1}(\tilde{w}^i)' \hat{\pi}_{J-1}(\tilde{w}^i) = \left\| \hat{\pi}_{J-1}(\tilde{w}^i) \right\|^2.$$

Again recall that $\hat{\pi}_{J-1}(\tilde{w}^i)$ is the vector of values of polynomials $\hat{P}_n(w)$, $n = 0, \dots, J-1$, evaluated at the given root \tilde{w}^i of the polynomial $\hat{P}_J(w)$. It is also an eigenvector of the Jacobi matrix $\hat{\Lambda}$ associated with eigenvalue \tilde{w}^i . We can therefore construct the weights \tilde{f}^i from the eigenvectors $\hat{\pi}_{J-1}(\tilde{w}^i)$, $i = 1, \dots, J$, we only need to determine an appropriate normalization of the eigenvectors.

Denote ϕ^i an arbitrarily scaled eigenvector of $\hat{\Lambda}$ associated with eigenvalue \tilde{w}^i . This vector can differ from $\hat{\pi}_{J-1}(\tilde{w}^i)$ only by a scaling factor. Since the first element of $\hat{\pi}_{J-1}(\tilde{w}^i)$ is equal to $\hat{P}_0(\tilde{w}^i) = \mu_0^{-1/2}$, we must have

$$\frac{\phi^i}{\phi_1^i} = \frac{\hat{\pi}_{J-1}(\tilde{w}^i)}{\hat{P}_0(\tilde{w}^i)}$$

have

$$\phi^i = \frac{\phi_1^i}{\hat{P}_0(\tilde{w}^i)} \hat{\pi}_{J-1}(\tilde{w}^i) = \sqrt{\mu_0} \phi_1^i \hat{\pi}_{J-1}(\tilde{w}^i).$$

As a consequence, the weights \tilde{f}^i are given by

$$\tilde{f}^i = \frac{1}{\left\| \hat{\pi}_{J-1}(\tilde{w}^i) \right\|^2} = \frac{\mu_0 (\phi_1^i)^2}{\left\| \phi^i \right\|^2}. \quad (2.36)$$

These results give us an algorithm for the construction of the quadrature nodes and weights:

1. Choose an interval (\underline{w}, \bar{w}) and a weighting function $f(w)$. Pick J .
2. Construct the monic polynomial base $P_n(w)$, $n = 0, \dots, J$ using the recurrence equa-

tion (2.32), respectively its orthonormal version $\hat{P}_n(w)$ using (2.33).

3. Construct the matrix $\hat{\Lambda}$ in (2.35) from the coefficients of the polynomials.
4. Find the eigenvalue \bar{w}^i of $\hat{\Lambda}$ and associated eigenvector ϕ^i , $i = 1, \dots, J$. Construct the weights \tilde{f}^i using the normalization (2.36) with

$$\mu_0 = \int_{\underline{w}}^{\bar{w}} f(w) dw = \left(\hat{P}_0(\bar{w}^i) \right)^{-2}.$$

The construction of the polynomial base requires the recursive computation of coefficients $a_{r,s}$ in (2.31). In general, this may require a numerical evaluation but well-known analytical formulas exist for special cases. In what follows, we derive specific formulas for particular weighting functions and integration ranges.

Gauss–Legendre quadrature

The Gauss–Legendre rule concerns the integration problem with $f(w) = 1$ and $(\underline{w}, \bar{w}) = (-1, 1)$. In this case, the orthogonal basis consists of the so-called Legendre polynomials. These satisfy the recursion

$$(n+1)\tilde{P}_{n+1}(w) = (2n+1)w\tilde{P}_n(w) - n\tilde{P}_{n-1}(w)$$

with $\tilde{P}_0(w) = 1$ and $\tilde{P}_1(w) = w$. These polynomials satisfy

$$\int_{-1}^1 \tilde{P}_n(w) \tilde{P}_m(w) f(w) dw = \frac{2}{2n+1} \delta_{mn},$$

so they are neither orthonormal nor monic. We want their monic representation in the form of the recursive formula (2.32), so that we can construct the matrix $\hat{\Lambda}$ in (2.35) from coefficients $a_{n,n}$ and $a_{n,n-1}$ of the monic polynomials. Therefore define the scaling factor λ_n as

$$P_n(w) = \lambda_n \tilde{P}_n(w)$$

where $P_n(w)$ is monic. Then we can rewrite the recurrence formula as

$$\begin{aligned} (n+1) \frac{P_{n+1}(w)}{\lambda_{n+1}} &= (2n+1)w \frac{P_n(w)}{\lambda_n} - n \frac{P_{n-1}(w)}{\lambda_{n-1}} \\ P_{n+1}(w) &= \frac{\lambda_{n+1}}{\lambda_n} \frac{2n+1}{n+1} w P_n(w) - \frac{\lambda_{n+1}}{\lambda_{n-1}} \frac{n}{n+1} P_{n-1}(w). \end{aligned}$$

The first two polynomials are still $P_0(w) = 1$ and $P_1(w) = w$. If $P_n(w)$ is monic then in order for $P_{n+1}(w)$ also to be monic, we need

$$\frac{\lambda_{n+1}}{\lambda_n} \frac{2n+1}{n+1} = 1 \quad \implies \quad \frac{\lambda_{n+1}}{\lambda_n} = \frac{n+1}{2n+1}.$$

This also implies that

$$\frac{\lambda_{n+1}}{\lambda_{n-1}} = \frac{\lambda_{n+1}}{\lambda_n} \frac{\lambda_n}{\lambda_{n-1}} = \frac{n+1}{2n+1} \frac{n}{2n-1} = \frac{n^2+n}{4n^2-1},$$

so that the recurrence formula becomes

$$P_{n+1}(w) = wP_n(w) - \frac{\lambda_{n+1}}{\lambda_{n-1}} \frac{n}{n+1} P_{n-1}(w) = wP_n(w) - \frac{n^2}{4n^2-1} P_{n-1}(w).$$

Hence

$$a_{n,n} = 0 \quad a_{n,n-1} = \frac{n^2}{4n^2-1},$$

which yields the $J \times J$ matrix $\hat{\Lambda}$ with a zero main diagonal and terms $\sqrt{a_{n,n-1}}$, $n = 1, \dots, J-1$ on the super- and subdiagonal. The eigenvalues of $\hat{\Lambda}$ correspond to the nodes \tilde{w}^j and the weights \tilde{f}^j can be computed from the associated eigenvectors.

The nodes and weights can also be directly translated to an approximation over an arbitrary interval (\underline{w}, \bar{w}) and weighting function $f(w) = c$ through a simple linear transformation. In this case, the new nodes w^j and weights f^j are related to \tilde{w}^j and \tilde{f}^j through

$$\begin{aligned} w^j &= \underline{w} + \frac{\bar{w} - \underline{w}}{2} (\tilde{w}^j + 1) \\ f^j &= c \frac{\bar{w} - \underline{w}}{2} \tilde{f}^j \end{aligned}$$

For example, when $f(w)$ is a density on (\underline{w}, \bar{w}) , then $c = (\bar{w} - \underline{w})^{-1}$

Gauss–Hermite quadrature

The Gauss–Hermite rule involves the weighting function $f(w) = \exp(-w^2)$ on the interval $(\underline{w}, \bar{w}) = (-\infty, +\infty)$. This weighting function satisfies

$$\int_{-\infty}^{\infty} f(w) dw = \sqrt{\pi}.$$

The orthogonal basis is given by the Hermite polynomials $\tilde{P}_n(w)$ given by $\tilde{P}_0(w) = 1$, $\tilde{P}_1(w) = 2w$, and further by the recurrence relation

$$\tilde{P}_{n+1}(w) = 2w\tilde{P}_n(w) - 2n\tilde{P}_{n-1}(w).$$

These polynomials satisfy the orthogonality restriction

$$\int_{-\infty}^{\infty} \tilde{P}_m(w) \tilde{P}_n(w) \exp(-w^2) dw = \sqrt{\pi} 2^n n! \delta_{mn}.$$

To construct the monic polynomials $P_n(w)$, we again define

$$P_n(w) = \lambda_n \tilde{P}_n(w).$$

Then we can rewrite the recurrence formula as

$$P_{n+1}(w) = 2 \frac{\lambda_{n+1}}{\lambda_n} w P_n(w) - 2n \frac{\lambda_{n+1}}{\lambda_{n-1}} P_{n-1}(w).$$

The first two polynomials are still $P_0(w) = 1$ and $P_1(w) = w$. If $P_n(w)$ is monic then in order for $P_{n+1}(w)$ also to be monic, we need

$$\frac{\lambda_{n+1}}{\lambda_n} = \frac{1}{2}.$$

This also implies that

$$\frac{\lambda_{n+1}}{\lambda_{n-1}} = \frac{\lambda_{n+1}}{\lambda_n} \frac{\lambda_n}{\lambda_{n-1}} = \frac{1}{4},$$

so that the recurrence formula becomes

$$P_{n+1}(w) = w P_n(w) - \frac{n}{2} P_{n-1}(w).$$

Hence the coefficients $a_{n,n}$ and $a_{n,n-1}$ in the general formula (2.32) that enter the tridiagonal matrix $\hat{\Lambda}$ in (2.35) satisfy

$$a_{n,n} = 0 \quad a_{n,n-1} = \frac{n}{2}.$$

The $J \times J$ matrix $\hat{\Lambda}$ therefore has zero main diagonal terms and terms $\sqrt{a_{n,n-1}} = \frac{n}{2}$, $n = 1, \dots, J-1$ on the super- and subdiagonal. The eigenvalues of $\hat{\Lambda}$ correspond to the nodes \tilde{w}^j and the weights \tilde{f}^j can be computed from the associated eigenvectors using formula (2.36).

The Gauss-Hermite rule is frequently used to compute expectations of a normally distributed random variable with mean μ and variance σ^2 . The weighting function is then the density

$$f(w) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2} \frac{(w-\mu)^2}{\sigma^2}\right).$$

Since the density integrates to one, and the new variable is a linear transformation of the original one, the nodes and weights are transformed as

$$w^j = \mu + \sqrt{2\sigma} \tilde{w}^j \quad f^j = \frac{1}{\sqrt{\pi}} \tilde{f}^j, \quad j = 1, \dots, J.$$

Chapter 3

Perturbation methods and linear state space models

Textbook: [Ljungqvist and Sargent \(2018\)](#), Chapter 2 (Sections 2.4–2.6, linear vector autoregressions, Sections 2.12–2.13, applications), Chapter 5 (linear quadratic dynamic programming). [Judd \(1998\)](#), Chapter 13 (perturbation methods). [Holmes \(1995\)](#) (series expansion methods).

Perturbation methods: [Blanchard and Kahn \(1980\)](#), [Sims \(2002\)](#), [Lombardo \(2010\)](#), [Borovička and Hansen \(2014\)](#), [Bhandari et al. \(2019\)](#), [Campbell and Shiller \(1988\)](#).

QuantEcon: [Quantitative Economics with Python](#): Topic 24 (AR(1) processes), Topic 27–28 (linear state space models), Topic 78 (linear regression). [Advanced Quantitative Economics with Python](#): Topics 16–19 (dynamic linear economies).

In this chapter, we study methods for finding approximate solutions to models that are potentially highdimensional, such as those used in quantitative dynamic macroeconomics. To achieve this, we need to find a class of stochastic processes that allow for tractable solutions of equilibrium models consisting of forward- and backward-looking equations, as well as their estimation. For this purpose, we study in Section 3.1 multivariate linear models, so-called linear vector autoregressions.

As a second step, we need to develop a numerical approach how to approximate interesting nonlinear models using linear dynamics. Here, we rely on perturbation methods that linearize the model in the vicinity of the model's steady state. The specific conceptual approach we use is the series expansion method.

The perturbation approximation is highly scalable and easy to implement but requires the model to be sufficiently smooth in the neighborhood of the steady state. It is necessary to be aware of the limitations of the method, in particular what features of the model will be neglected by the approximation. A linear approximation of the model will intuitively work well if nonlinearities are not substantial. Higher-order polynomial approximations can deal with some nonlinearities as well, and standard packages such as Dynare can construct such polynomial approximations of the model solution automatically. However,

evaluation of the accuracy of the approximation is a nontrivial issue, as mentioned in the cautious note in Section 3.2.8 at the end of this chapter.

3.1 Linear vector autoregressions

We will now study a class of multivariate linear models. These models are easy to work with and relatively easy to estimate using well-understood methods, and can often be the outcome of an approximation to a solution of a nonlinear equilibrium model.

Analogously to the Markov chain model described in Section 1.1, the **multivariate stochastic linear model** is described by the following components:

- state of the system $x_t \in \mathbb{R}^n$,
- initial distribution $\pi_0(x_0) \sim N(\mu_0, \Sigma_0)$,
- transition density $\pi(x' | x) \sim N(A_0 x, CC')$, where A_0 is an $n \times n$ matrix and C is an $n \times p$ matrix.

Observe that this specification again describes completely a filtered probability space from Definition A.3. The **sample space** Ω contains all infinite sequences of realizations $x^\infty = (x_0, x_1, \dots)$. The **filtration** $\{\mathcal{F}_t\}_{t=0}^\infty$ is generated from measurable sets of partial histories $x^t = (x_0, x_1, \dots, x_t)$. The construction is somewhat more involved here since the state space is infinite. In particular, for a given t , the σ -algebra \mathcal{F}_t is constructed from all sets of paths $x^\infty \in \Omega$ that take the following form

$$F_t = \{x^\infty \in \Omega : x_0 \in B_0, x_1 \in B_1, \dots, x_t \in B_t\}$$

where B_j , $j = 0, \dots, t$ are open (or Borel) sets on \mathbb{R}^n , and then from all sets that can be constructed by repeated application of the rules in Definition A.1. The σ -algebra \mathcal{F} can then be defined formally as $\mathcal{F} = \bigcup_{t \geq 0} \mathcal{F}_t$. Finally, the probability measure P describing the full joint distribution of the model can be formed from the initial distribution and the transition density.

It is extremely useful to notice that the model can be equivalently represented using the **stochastic linear difference equation**

$$x_{t+1} = A_0 x_t + C w_{t+1} \quad w_{t+1} \sim N(0, I_p) \quad \text{iid.} \quad (3.1)$$

where w_{t+1} is an $p \times 1$ vector of **iid Gaussian** shocks (so-called random innovations). Such a model is called a **vector autoregression** (VAR). Observe that when this equation can be inverted to obtain

$$w_{t+1} = C^{-1}(x_{t+1} - A_0 x_t)$$

then the information sets \mathcal{F}_t can be equivalently generated using partial histories $x^t = (x_0, x_1, \dots, x_t)$ or using the histories of innovations (x_0, w_1, \dots, w_t) .

Some of the results that follow will continue to hold even when the assumptions on the distribution of innovations are weakened. In particular, we can relax the assumption of a Gaussian distribution, and instead assume that w_{t+1} is a random vector satisfying

$$\begin{aligned} E[w_{t+1} | \mathcal{F}_t] &= 0 \\ E[w_{t+1}w'_{t+1} | \mathcal{F}_t] &= I_p, \end{aligned} \quad (3.2)$$

where \mathcal{F}_t is the σ -algebra (information set) generated by (x_0, w_1, \dots, w_t) . The sequence of shocks $\{w_{t+1}\}_{t=0}^{\infty}$ satisfying (3.2) is called a **martingale difference sequence**. An even weaker assumption further relaxes the conditional moments, and only assumes that the shocks are unconditionally mean zero and uncorrelated

$$\begin{aligned} E[w_{t+1}] &= 0 \\ E[w_t w'_{t-j}] &= I_p \cdot \mathbf{1}\{j=0\}. \end{aligned} \quad (3.3)$$

A sequence of shocks satisfying the pair of restrictions (3.3) is called **white noise**.

We will often append an **observation equation**, or **measurement equation**, to obtain what is called a **state-space representation** of the model:

$$\begin{aligned} x_{t+1} &= A_o x_t + C w_{t+1} \\ y_t &= G x_t + v_t \end{aligned} \quad (3.4)$$

where y_t and v_t are $m \times 1$ vectors. The vector y_t represents observations of a potentially 'hidden' state x_t , and v_t is iid measurement noise with a given covariance matrix. In some applications, we will allow v_t to be correlated with w_t .

The multivariate linear Markov process is incredibly versatile because many interesting linear models can be rewritten into the **VAR form**.

Example 3.1. *A scalar second-order autoregression*

$$z_{t+1} = \alpha + \rho_1 z_t + \rho_2 z_{t-1} + w_{t+1} \quad (3.5)$$

can be written as

$$x_{t+1} = \begin{bmatrix} z_{t+1} \\ z_t \\ 1 \end{bmatrix} = \begin{bmatrix} \rho_1 & \rho_2 & \alpha \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} z_t \\ z_{t-1} \\ 1 \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} w_{t+1}$$

with measurement equation

$$z_t = [1 \ 0 \ 0] x_t$$

The matrix A_o constructed above is called the *companion form*.

The autoregressive model (3.5) is an example of how uncorrelated disturbances w_t may generate persistent oscillations in the observed series z_t . This effect was independently described by Eugen Slutsky and Udney Yule in [Slutsky \(1927\)](#) (appeared in English as [Slutsky](#)

(1937)) and Yule (1927), and is known as the **Slutsky–Yule effect**. Slutsky (1927) noted, for example, that moving averages constructed from the random numbers drawn in the Russian government lottery resemble the time series of British business cycles.

The idea revolutionized the way how to think about the propagation mechanism generating business cycles. Ragnar Frisch constructed a continuous-time model of aggregate dynamics in Frisch (1933) in which he distinguishes between the ‘**impulse problem**’ and the ‘**propagation problem**’. Oscillations in his model are generated by a time-to-build mechanism where capital goods need time to be completed before they can be used for production, an early precursor to the time-to-build model of Kydland and Prescott (1982). As another early example, the equilibrium dynamics for aggregate output in the **multiplier-accelerator model** of Samuelson (1939) take exactly the form (3.5), where ρ_1 and ρ_2 are model parameters calibrated to mimic the characteristics of business cycle fluctuations, generated by fluctuations in government spending.

Example 3.2. *The vector autoregression can incorporate moving-average dynamics by stacking the history of shocks. The ARMA(1,1) model*

$$z_{t+1} = \rho z_t + w_{t+1} + \gamma w_t$$

can be written as

$$x_{t+1} = \begin{bmatrix} z_{t+1} \\ w_{t+1} \end{bmatrix} = \begin{bmatrix} \rho & \gamma \\ 0 & 0 \end{bmatrix} \begin{bmatrix} z_t \\ w_t \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} w_{t+1}$$

with measurement equation

$$z_t = [1 \ 0] x_t.$$

Other examples of models that can be suitably stacked into the VAR form include an **order- k vector autoregression**

$$z_{t+1} = \sum_{j=1}^k A_j z_{t+1-j} + C_y w_{t+1},$$

or models that include deterministic or stochastic **seasonality**:

$$\begin{aligned} y_t &= y_{t-4} \\ y_t &= \phi y_{t-4} + w_t. \end{aligned}$$

3.1.1 First and second moments

When innovations w_{t+1} are normally distributed and the unconditional distribution of the initial state x_0 is normal as well, then the linear form of (3.1) implies that x_t will be normally distributed as well. Since normal distributions are completely described by their first two moments, tracing the first two moments over time is sufficient for the description of the joint distribution of the process.

However, even when the innovations are not normal, we are often still interested in studying the dynamics of the first two moments. This leads us to the definition of covariance stationarity.

Definition 3.1. A stochastic process is said to be *covariance stationary* if

- the mean is independent of time, $E[x_t] = E[x_0] = \bar{\mu}$
- the sequence of autocovariance matrices

$$E \left[(x_t - E[x_t]) (x_{t+j} - E[x_{t+j}])' \right]$$

only depends on j , not on t .

Obviously, a stationary process (according to Definition A.11) is covariance stationary. On the other hand, a linear covariance stationary process with normal innovations and normal unconditional distribution of the initial state is also stationary.

Definition 3.2. A real square matrix A_0 is said to be *stable* if all its eigenvalues are strictly within the unit circle.

In order for x_{t+1} to have a stationary mean different from zero, it will often be useful to impose a particular structure on

$$x_{t+1} = A_0 x_t + C w_{t+1} \quad (3.6)$$

by singling out a constant from the evolution of the state:

$$\begin{bmatrix} x_{1,t+1} \\ x_{2,t+1} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ b & \tilde{A} \end{bmatrix} \begin{bmatrix} x_{1,t} \\ x_{2,t} \end{bmatrix} + \begin{bmatrix} 0 \\ \tilde{C} \end{bmatrix} w_{t+1} \quad x_0 = \begin{bmatrix} x_{1,0} \\ x_{2,0} \end{bmatrix} \quad (3.7)$$

where $x_{1,t}$ is scalar. The matrix A_0 then has one unit root, and the remaining roots are the roots of \tilde{A} which we assume is stable.

Law of motion for first moments

Denote $\mu_t \doteq E[x_t]$ the unconditional mean of x_t (which can still depend on t if the process is not stationary). Then

$$\mu_{t+1} = A_0 \mu_t$$

and we can find $\lim_{t \rightarrow \infty} \mu_t = \mu_\infty$ as the unique solution to

$$\bar{\mu} = A_0 \bar{\mu} \quad \implies \quad (I - A_0) \bar{\mu} = 0$$

To provide more information let us look at the structured equation (3.7). Denote $\bar{\mu}' = (\bar{\mu}_1, \bar{\mu}'_2)$. Obviously, $x_{1,t} = x_{1,0} = \bar{\mu}_1$. The lower block can be written as

$$x_{2,t+1} = \bar{a} x_{1,0} + \tilde{A} x_{2,t} + \tilde{C} w_{t+1}.$$

Taking unconditional expectations, we have

$$\mu_{2,t+1} = b \bar{\mu}_1 + \tilde{A} \mu_{2,t}.$$

When \tilde{A} is stable, $\lim_{t \rightarrow \infty} \mu_{2,t} = \mu_{2,\infty} = \bar{\mu}_2$, and we can find the value as

$$\bar{\mu}_2 = (I - \tilde{A})^{-1} b \bar{\mu}_1.$$

Law of motion for second moments

In order to derive the evolution of unconditional variance, denote

$$\Sigma_t \doteq E [(x_t - \mu_t) (x_t - \mu_t)'].$$

The law of motion for Σ_t can be derived from (3.6) by subtracting the unconditional mean $\mu_{t+1} = A_0 \mu_t$ from both sides and taking the variance of both sides. Hence

$$\Sigma_{t+1} = A_0 \Sigma_t A_0' + C C'.$$

A fixed point of this recursion satisfies

$$\Sigma_\infty = A_0 \Sigma_\infty A_0' + C C'. \quad (3.8)$$

We will denote this fixed point $C_x(0) = \Sigma_\infty$. This fixed point is the covariance matrix

$$C_x(0) = E [(x_t - \bar{\mu}) (x_t - \bar{\mu})']$$

under the stationary distribution. Equation (3.8) is a **discrete Lyapunov equation** and can be efficiently solved using alternative algorithms (like the **doubling algorithm**).

Similarly, to compute the autocovariance function $C_x(j)$, start with (3.6) and write

$$\begin{aligned} x_{t+j} - \mu_{t+j} &= A_0 (x_{t+j-1} - \mu_{t+j-1}) + C w_{t+j} = \dots \\ &= A_0^j (x_t - \mu_t) + A_0^{j-1} C w_{t+1} + \dots + C w_{t+j} \end{aligned}$$

Post-multiply by $(x_t - \mu_t)'$ and take unconditional expectations to obtain

$$E [(x_{t+j} - \mu_{t+j}) (x_t - \mu_t)'] = A_0^j E [(x_t - \mu_t) (x_t - \mu_t)']$$

Hence, when the process has a stationary mean, we obtain

$$C_x(j) = A_0^j C_x(0).$$

The sequence $\{C_x(j)\}_{j=0}^\infty$ is the autocovariance function or autocovariogram.

Example 3.3. Consider $y_t = G x_t$. Then $C_y(j) = G C_x(j) G'$.

Summary

To summarize, we distinguish different moments based on the conditioning we impose:

- conditional moments $E[x_{t+1} | x_t] = Ax_t$, $Cov(x_{t+1} | x_t) = CC'$
- moments conditional on x_0 ,

$$E[x_t | x_0] = E_0[x_t] = A_o^t x_0$$

$$E[(x_t - E_0[x_t])(x_t - E_0[x_t])'] = \sum_{h=0}^{t-1} A_o^h CC' (A_o^h)'$$

- unconditional moments $E[x_t] \doteq \mu_t$ and $E[(x_t - \mu_t)(x_t - \mu_t)'] = \Sigma_t$, satisfying

$$\mu_{t+1} = A_o \mu_t$$

$$\Sigma_{t+1} = A_o \Sigma_t A_o' + CC'$$

- stationary moments

$$(I - A_o) \bar{\mu} = 0$$

$$C_x(0) = A_o C_x(0) A_o' + CC'$$

$$C_x(j) = A_o^j C_x(0)$$

3.1.2 Impulse response functions

The stochastic process we posited in (3.1) specifies a law of motion that describes a deterministic propagation mechanism for x_t , systematically perturbed by random innovations w_{t+1} . This idea goes back to the impulse and propagation problems described by [Frisch \(1933\)](#):

“There are several alternative ways in which one may approach the impulse problem. . . One way which I believe is particularly fruitful and promising is to study what would become of the solution of a determinate dynamic system if it were **exposed to a stream of erratic shocks** that constantly upsets the continuous evolution, and by so doing introduces into the system the energy necessary to maintain the swings.”

In order to understand the propagation mechanism, we want to capture how a shock today affects the distribution of the stochastic process in the future. In general, the system starts in a state x_0 and is impacted in periods $t = 1, 2, \dots$ by the sequence of shocks w_t . Specifically, we want to pose the following question: What are the consequences of perturbing the distribution of the shock w_1 today for the distribution of x_t , $t \geq 1$?

In order to answer it, consider a common initial condition x_0 and two alternative processes representing iid disturbances:

$$w = \{w_1, w_2, w_3, \dots\}$$

$$\tilde{w} = \{\tilde{w}_1, \tilde{w}_2, \tilde{w}_3, \dots\}$$

where we impose that $\tilde{w}_j = w_j, \forall j \geq 2$. The innovation processes thus have the same innovations except period one. Now define

$$\begin{aligned} x_{t+1} &= A_o x_t + C w_{t+1}, & t \geq 0 \\ \tilde{x}_{t+1} &= A_o \tilde{x}_t + C \tilde{w}_{t+1} \end{aligned}$$

The **impulse response function** is defined as the difference

$$\tilde{x}_{t+1} - x_{t+1}.$$

Observe that the impulse response function is a stochastic process that in general depends on x_0 , as well as on both sequences of shocks \tilde{w} and w . So in the general nonlinear case, we need to think about ways how to summarize information contained in the process $\tilde{x}_{t+1} - x_{t+1}$. It turns out that in the **linear case**, this difference takes a convenient simple form.

In order to derive it, notice that

$$\begin{aligned} x_{t+1} &= A_o x_t + C w_{t+1} = A_o^2 x_{t-1} + A_o C w_t + C w_{t+1} \\ &= A_o^{t+1} x_0 + \sum_{j=0}^t A_o^j C w_{t+1-j} \end{aligned} \quad (3.9)$$

This is the so-called **moving-average representation** of the process $\{x_t\}$ that specifies the process as a linear combination of past innovations. Using this moving-average representation, it is easy to infer that

$$\tilde{x}_{t+1} - x_{t+1} = A_o^t C (\tilde{w}_1 - w_1)$$

In the study of linear models, it is a common choice to take $w_1 = 0$ and $\tilde{w}_1 = e_k, k = 1, \dots, p$. The matrix-valued function

$$h_t = A_o^t C \quad (3.10)$$

is therefore also commonly referred to as (linear) **impulse response function**.

Observe that h_t is a matrix whose entry $[A_o^t C]_{ik}$ represents the response t periods after the impact of the shock of the i -th element of the vector state process x to a perturbation of the k -th element of the innovation w_1 . Finally, notice that the representation (3.9) makes the computation of the conditional covariance of x_t easy, because w are iid innovations. Hence

$$E \left[(x_t - E[x_t | x_0]) (x_t - E[x_t | x_0])' \right] = \sum_{j=0}^{t-1} A_o^j C C' (A_o^j)'$$

3.1.3 Applications

Prediction and discounting

The moving average representation is an extremely powerful tool for computing expectations and other statistics. Rewrite (3.9) as

$$x_{t+j} = A_o^j x_t + \sum_{k=0}^{j-1} A_o^k C w_{t+j-k}$$

Hence

$$E_t [x_{t+j}] \doteq E [x_{t+j} | x_t] = A_o^j x_t.$$

Similarly, consider a function $y_t = Gx_t$ where G is a conformable matrix. Then

$$E_t \left[\sum_{j=0}^{\infty} \beta^j y_{t+j} \right] = G \sum_{j=0}^{\infty} (\beta A_o)^j x_t = G (I - \beta A_o)^{-1} x_t$$

provided that the matrix βA_o has all unit roots smaller than one in modulus.

Geometric sums of quadratic forms

In linear-quadratic models, we often want to calculate

$$\alpha_t = E_t \left[\sum_{j=0}^{\infty} \beta^j x'_{t+j} Y x_{t+j} \right]. \quad (3.11)$$

We can proceed by guess and verify and establish a recursive formula

$$\begin{aligned} \alpha_t &= x'_t Y x_t + \beta E_t \left[E_{t+1} \sum_{j=0}^{\infty} \beta^j x'_{t+1+j} Y x_{t+1+j} \right] = \\ &= x'_t Y x_t + \beta E_t \alpha_{t+1}. \end{aligned}$$

Guessing

$$\alpha_t = x'_t v x_t + \sigma$$

where v is a symmetric $n \times n$ matrix and σ a scalar, we plug in to obtain

$$\begin{aligned} x'_t v x_t + \sigma &= x'_t Y x_t + \beta E_t [x'_{t+1} v x_{t+1} + \sigma] \\ &= x'_t Y x_t + \beta x'_t A'_o v A_o x_t + \beta E_t [w'_{t+1} C' v C w_{t+1}] + \beta \sigma \\ &= x'_t Y x_t + \beta x'_t A'_o v A_o x_t + \beta \text{tr} [C' v C] + \beta \sigma. \end{aligned}$$

Comparing coefficients on constant terms and terms involving squares of x_t , we obtain

$$\begin{aligned} \nu &= Y + \beta A'_o \nu A_o \\ \sigma &= (1 - \beta)^{-1} \beta \text{tr} [C' \nu C] \end{aligned} \quad (3.12)$$

where $\text{tr} [\cdot]$ denotes the trace of a matrix. The equation for ν is again a **discrete Lyapunov equation**.

Asset pricing

An area that heavily relies on linear methods (in a specific sense) is asset pricing. Generally, the current ‘fundamental’ price of an asset is the sum of the present discounted value of future cash flows that the asset pays out.

We therefore need to determine a model of cash flows, and a model of discounting. This is the content, in an abstract general equilibrium setup, of the Theory of Value of [Debreu \(1959\)](#). A model of horizon-dependent discounting can be traced back to the study of the yield curve in [Fisher \(1896, 1906, 1930\)](#), and [Hicks \(1939\)](#). [Arrow \(1964\)](#) studied state-dependent discounting in a model of risky payoffs, leading to so-called **Arrow–Debreu prices**, which are prices of hypothetical securities with a unit payoff in a particular state of the economy.

Consider the **cash flow** y_t and the ‘discount factor’ z_t , modeled as

$$y_t = Gx_t \quad z_t = Hx_t$$

where G and H are row vectors. We are interested in computing the ‘asset price’

$$p_t = E_t \left[\sum_{j=0}^{\infty} \beta^j z_{t+j} y_{t+j} \right]$$

Observe that given the **stochastic discount factor** process $\beta^j z_{t+j}$, the asset price is a linear function of the cash flows. The stochastic discount factor is typically derived from agent’s preferences, reflecting the marginal rate of substitution between today and uncertain future states. We derived the stochastic discount factor in this way in Section 1.2. Moreover,

$$p_t = E_t \left[\sum_{j=0}^{\infty} \beta^j x'_{t+j} H' G x_{t+j} \right],$$

and we can use (3.11) to evaluate this sum to obtain

$$p_t = x'_t \nu x_t + \sigma.$$

The coefficients ν and σ are determined in (3.12). Recall that σ is determined as the discounted sum of covariances of the innovations in z_t and y_t . Specifically, observe that the term $C' \nu C$ depends on the underlying volatility of x_t reflected in C , as well as on $Y = H' G$

that reflects the comovement of cash flows with the stochastic discount factor. The value of σ thus can be interpreted as a **risk premium** on asset with cash flow y_t .

Evaluation of a dynamic criterion

We now consider a recursive computational algorithm that provides the basis for linear-quadratic dynamic programming. Let the system be driven by

$$x_{t+1} = Ax_t + Bu_t + Cw_{t+1}$$

where u_t is a control vector chosen by the decision maker as

$$u_t = -Fx_t.$$

Later, we will determine F as an outcome of an optimization problem. For the given control u_t , we want to compute

$$v(x_0; F) = -E_0 \left[\sum_{t=0}^{\infty} \beta^t (x_t' R x_t + u_t' Q u_t) \right],$$

where R and Q are symmetric matrices. We can rewrite this problem as

$$\begin{aligned} x_{t+1} &= A_o x_t + Cw_{t+1} \\ v(x_0; F) &= -E_0 \left[\sum_{t=0}^{\infty} \beta^t x_t' Y x_t \right] \end{aligned}$$

with

$$A_o = A - BF \quad Y = R + F' Q F.$$

The solution is then given by

$$v(x_0; F) = -x_0' P x_0 - \sigma$$

where, in line with equations (3.12), P satisfies

$$\begin{aligned} P &= Y + \beta A_o' P A_o \\ &= R + F' Q F + \beta (A - BF)' P (A - BF) \end{aligned} \tag{3.13}$$

Finding optimal control

Imagine we are now interested in finding the optimal control $\{u_t\}_{t=0}^{\infty}$ that maximizes $v(x_0)$:

$$v(x_0) = \max_{\{u_t\}_{t=0}^{\infty}} -E_0 \left[\sum_{t=0}^{\infty} \beta^t (x_t' R x_t + u_t' Q u_t) \right].$$

We describe here a variant of the value function iteration algorithm called **Howard improvement algorithm**, described in Howard (1960), Chapter 4. This algorithm is also called the **policy function iteration**, see Chapter 5 in Ljungqvist and Sargent (2018)).

Conjecture an initial guess for the policy function in the form $u_t = -F_0x_t$ for some matrix F_0 . Let us assume that the decision maker uses F_0 from period 1 onward. Then

$$v(x_1; F_0) = -x_1' P_0 x_1 - \sigma_0,$$

where P_0 is given as the solution to (3.13) for $F = F_0$, and σ_0 is a constant that is inconsequential for the following arguments. Given this restriction on future controls, let us choose the optimal control at time 0 by solving

$$\max_{u_0} - \{x_0' R x_0 + u_0' Q u_0 + \beta E_0 [\underbrace{(Ax_0 + Bu_0 + Cw_1)' P_0 (Ax_0 + Bu_0 + Cw_1)}_{x_1' P_0 x_1}] + \beta \sigma_0\}$$

We now differentiate with respect to the vector u_0 . Notice that for a function $f(u) : \mathbb{R}^k \rightarrow \mathbb{R}$, the first derivative is the column vector

$$\frac{\partial}{\partial u} f(u) = \left(\frac{\partial}{\partial u_1} f(u), \dots, \frac{\partial}{\partial u_k} f(u) \right)'$$

so you can verify that, for example,

$$\frac{\partial}{\partial u} (u' Q u) = Qu + Q'u,$$

which, for a symmetric matrix Q , reduces to $\frac{\partial}{\partial u} (u' Q u) = 2Qu$. The first-order condition implies

$$0 = -2Qu_0 - 2\beta B' P_0 (Ax_0 + Bu_0),$$

and hence

$$u_0 = -\underbrace{\beta (Q + \beta B' P_0 B)^{-1} B' P_0 A}_{\doteq F_1} x_0. \quad (3.14)$$

We can now construct the iterative step. Given F_j , solve equation (3.13) for P_j , and then improve the policy to F_{j+1} using equation (3.14):

$$\begin{aligned} P_j &= R + F_j' Q F_j + \beta (A - B F_j)' P_j (A - B F_j) \\ F_{j+1} &= \beta (Q + \beta B' P_j B)^{-1} B' P_j A. \end{aligned} \quad (3.15)$$

Initial F_0 should be chosen such that $\sqrt{\beta} (A - B F_0)$ has stable roots, in order for P_0 to be finite. When the process converged, $F_j \rightarrow F_\infty$, there is no one-period improvement in policy $u_t = -F_\infty x_t$. The algorithm also proves that optimal policy is linear in x_t .

Finally, observe that the matrices P_j and F_j do not depend on C . Hence optimal control $u_t = -F_\infty x_t$ does not depend on the volatility of the time series. This is a result known as

certainty equivalence—in linear-quadratic models, the presence of uncertainty does not alter decision rules. However, uncertainty affects the constant term in the value function, as we have seen in the asset pricing application.

Value function iteration and policy function iteration

An alternative to iterative scheme (3.15) is to iterate on both equations simultaneously:

$$\begin{aligned} P_{j+1} &= R + F_j' Q F_j + \beta (A - B F_j)' P_j (A - B F_j) \\ F_j &= \beta (Q + \beta B' P_j B)^{-1} B' P_j A. \end{aligned} \quad (3.16)$$

This time, the algorithm can be initialized from an arbitrary initial positive-semidefinite matrix P_0 ($P_0 = 0$ is a popular choice). However, this algorithm is slower in practice than (3.15).

To understand why, observe that algorithm (3.16) corresponds to the linear-quadratic version of the **value function iteration** algorithm from Section 2. Given the iterate of the value function given by matrix P_j , a new current-period decision $u = -F_j x$ is derived, and a new update of the value function given by P_{j+1} is constructed, under the assumption that the new decision is followed in the current period only.

On the other hand, the policy function iteration algorithm (3.15) assumes that given P_j and the new policy $u = -F_{j+1} x$, the new iterate for the value function, given by P_{j+1} , is computed by assuming that this new policy $u = -F_{j+1} x$ is obeyed in all periods (a new fixed point for the value function with the given policy is found).

3.1.4 Likelihood of the linear vector autoregressive process

In Section 1.1.5, we constructed the likelihood of data generated by a Markov chain. Here, we look at the likelihood associated with the linear vector process with Gaussian innovations. We are given a sample $\hat{x}^T = \{\hat{x}_t\}_{t=0}^T$ generated from model (3.1)

$$x_{t+1} = A_o x_t + C w_{t+1} \quad w_{t+1} \sim N(0, I_n)$$

where we for convenience assume that CC' is invertible. The joint density of the distribution of the process can be written as

$$f(x^T) = f(x_T | x_{T-1}) \dots f(x_1 | x_0) f(x_0),$$

where we used the Markov property to construct the factorization. For the Gaussian distribution,

$$f(x_{t+1} | x_t) = \frac{1}{(2\pi)^{n/2} (\det(CC'))^{1/2}} \exp\left(-\frac{1}{2} (x_{t+1} - A_o x_t)' (CC')^{-1} (x_{t+1} - A_o x_t)\right).$$

We assume that there is a vector of unknown parameters θ such that matrices A_o and C (and perhaps also the distribution $f(x_0)$) are functions of θ . The goal is then to maximize the log-likelihood

$$\begin{aligned} \log L(\theta | \hat{x}^T) &= \log f(\hat{x}_0 | \theta) - \frac{T}{2} \log \det (C(\theta) C(\theta)') \\ &\quad - \frac{1}{2} \sum_{t=0}^{T-1} (\hat{x}_{t+1} - A_o(\theta) \hat{x}_t)' (C(\theta) C(\theta)')^{-1} (\hat{x}_{t+1} - A_o(\theta) \hat{x}_t). \end{aligned}$$

In Section 4.3.1 we revisit this problem in a more complicated environment where x_t is not observable and we only observe imperfect signals of x_t . The estimator will then also involve a filtering method to estimate the unobserved path of x_t .

3.2 Perturbation methods

Linear vector autoregressions are appealing due to their tractability, allowing for easy treatment of high-dimensional problems. However, most economic models do not adhere to such a linear form. At the same time, it turns out that a range of stochastic dynamic equilibrium models can be suitably approximated on the relevant part of the state space using linear dynamics.

The idea of a linear perturbation method is to start with a particular convenient point in the state space that will serve as an expansion point. A convenient choice is the deterministic steady state, a point to which the trajectory in a deterministic version of the model without any uncertainty would converge. Then the approximation technique introduces a ‘small’ amount of uncertainty in the form of a perturbation of the dynamics in the vicinity of the steady, and constructs a linear approximation to study the approximate ‘local’ behavior. Finally, we extend the approximate local behavior to the level of uncertainty in the original model.

We first provide a derivation of the perturbation approximation using the so-called series expansion method, and then discuss several applications. First, we study linear approximations of asset return dynamics due to [Campbell and Shiller \(1988\)](#), and then consider log-linear approximations of price-dividend ratios widely used in the asset pricing literature. We then show how to approximate a nonlinear optimal control problem using a linear-quadratic approximation with a quadratic objective function and linear law of motion of the state, and map it to the framework analyzed in Section 3.1.3. Finally, we briefly discuss approaches to solving linear stochastic models with backward- and forward-looking equations arising in general equilibrium dynamics.

3.2.1 Taylor’s theorem

The idea of the perturbation approximation extends the logic of Taylor’s theorem to dynamic environments. Consider a function $f(x) : \mathbb{R} \rightarrow \mathbb{R}$ and a particular point $\bar{x} \in \mathbb{R}$. Taylor’s theorem shows that if $f(x)$ is k -times differentiable at \bar{x} , then $f(x)$ can be written

using a polynomial expansion

$$f(x) = f(\bar{x}) + f'(\bar{x})(x - \bar{x}) + \frac{f''(\bar{x})}{2}(x - \bar{x})^2 + \dots + \frac{f^{(k)}(\bar{x})}{k!}(x - \bar{x})^k + h_k(x)(x - \bar{x})^k$$

where $h_k(x)$ is a remainder (error) function such that

$$\lim_{x \rightarrow \bar{x}} h_k(x) = 0.$$

This means that the approximation error in the vicinity of \bar{x} is of higher order than $(x - \bar{x})^k$. For the ‘first-order’ case ($k = 1$), we obtain the linear approximation

$$f(x) = f(\bar{x}) + f'(\bar{x})(x - \bar{x}) + h_1(x)(x - \bar{x}) = f(\bar{x}) + f'(\bar{x})(x - \bar{x}) + o(x - \bar{x})$$

where $o(\varepsilon)$ is the ‘little-o’ function such that

$$\lim_{\varepsilon \rightarrow 0} \frac{o(\varepsilon)}{\varepsilon} = 0.$$

An rigorous application of Taylor’s theorem in the context of the dynamic macroeconomic model is a nontrivial extension, since it involves a simultaneous approximation of the whole dynamic system. More discussion is provided in [Judd \(1998\)](#), Chapter 13.

3.2.2 Series expansion method

We base the perturbation approximation on the series expansion method, originally developed for approximation of differential equations ([Holmes \(1995\)](#)), and applied in the context of dynamic macroeconomic models by [Lombardo \(2010\)](#). More discussion, various alternative approaches, and applications in the context of macroeconomics and asset pricing are provided in [Borovička and Hansen \(2014\)](#) and [Bhandari et al. \(2019\)](#).

Let x_t be an n -dimensional Markov stochastic process of the form

$$x_{t+1} = \psi(x_t, w_{t+1}) \quad w_{t+1} \sim N(0, I_p)$$

with a given initial condition x_0 , where w_{t+1} is a p -dimensional standard normal shock, and $\psi : \mathbb{R}^n \times \mathbb{R}^p \rightarrow \mathbb{R}^n$ is assumed to be a sufficiently smooth function.

We now extend this model by consider a class of models indexed by a perturbation parameter q that scales the volatility of the shocks w_{t+1} :

$$x_{t+1}(q) = \psi(x_t(q), qw_{t+1}, q). \quad (3.17)$$

Each choice of q generates different dynamics of the innovations, and hence a different solution $x_t(q)$ for the paths of the process. The last argument in (3.17) allows for the function ψ to also explicitly depend on the perturbation parameter, which is a choice that allows for additional flexibility in designing the perturbation method.

Specifically, for $q = 0$, we obtain the deterministic model

$$x_{t+1}(0) = \psi(x_t(0), 0, 0),$$

while for $q = 1$, we recover the original model

$$x_{t+1}(1) = \psi(x_t(1), w_{t+1}, 1).$$

Assume that we can write the class of stochastic processes $x_t(q)$ indexed by the perturbation parameter in the form of a series expansion around $q = 0$:

$$x_t(q) \approx x_{0t} + qx_{1t} + \frac{q^2}{2}x_{2t} + \dots \quad (3.18)$$

The series expansion extends the idea of Taylor expansion to a stochastic dynamic environment. The processes x_{jt} , $j \geq 0$, can be viewed as derivatives of x_t with respect to the perturbation parameter.

We want to construct tractable forms of the laws of motion for the derivatives x_{jt} , and then reconstruct the original process $x_t(q)$ for $q = 1$ using (3.18). The laws of motion can be inferred by substituting (3.18) into (3.17)

$$x_{0t+1} + qx_{1t+1} + \frac{q^2}{2}x_{2t+1} + \dots = \psi\left(x_{0t} + qx_{1t} + \frac{q^2}{2}x_{2t} + \dots, qw_{t+1}, q\right), \quad (3.19)$$

differentiating j times and evaluating the derivatives at $q = 0$.

The zero-th order derivative is computed by evaluating (3.19) at $q = 0$:

$$x_{0t+1} = \psi(x_{0t}, 0, 0).$$

We assume that the deterministic dynamics in the process x_{0t} converges to a steady state \bar{x} , which is a fixed point of the equation

$$\bar{x} = \psi(\bar{x}, 0, 0).$$

Solving this nonlinear equation for this fixed point, we set $x_{0t} = \bar{x}$. Differentiating (3.19) with respect to q and evaluating at $q = 0$ then yields

$$x_{1t+1} = \psi_q + \psi_x x_{1t} + \psi_w w_{t+1} \quad (3.20)$$

where ψ_x , and ψ_w are the $n \times n$ and $n \times p$ partial derivative matrices of ψ with respect to x and w , respectively, evaluated at $(\bar{x}, 0, 0)$, and ψ_q is the partial derivative with respect to q :

$$\psi_q = \frac{\partial \psi}{\partial q}(\bar{x}, 0, 0) \quad \psi_x = \frac{\partial \psi}{\partial x}(\bar{x}, 0, 0) \quad \psi_w = \frac{\partial \psi}{\partial w}(\bar{x}, 0, 0).$$

The laws of motion for higher-order derivatives x_{2t} , x_{3t} , ... can be constructed by continuing the procedure, we stop here with the linear approximation. We can then combine the

derivatives to obtain an approximation of x_t via (3.18) evaluated at $q = 1$.

We now move on to provide several fruitful applications of the perturbation approximation.

3.2.3 Campbell–Shiller decomposition of returns

In this section, we analyze an application of the series expansion technique to linear approximation of asset returns, first used by [Campbell and Shiller \(1988\)](#). This approximation will lead to a highly tractable analysis of the sources of variation in price-dividend ratios.

We start with the definition of the return

$$R_{t+1} = \frac{Q_{t+1} + G_{t+1}}{Q_t} = \frac{Q_{t+1}/G_{t+1} + 1}{Q_t/G_t} \frac{G_{t+1}}{G_t}, \quad (3.21)$$

and rewrite it in logarithms of the returns, price-dividend ratio, and dividend growth

$$r_{t+1} = \log R_{t+1} \quad z_t = \log \frac{Q_t}{G_t} \quad g_{t+1} = \log \frac{G_{t+1}}{G_t}$$

to obtain

$$\exp(r_{t+1}) = \frac{\exp(z_{t+1}) + 1}{\exp(z_t)} \exp(g_{t+1}) \quad (3.22)$$

Now we apply the first-order series expansion to all three quantities:

$$\begin{aligned} r_{t+1} &\approx \bar{r} + qr_{1t+1} \\ z_t &\approx \bar{z} + qz_{1t} \\ g_{t+1} &\approx \bar{g} + qg_{1t+1} \end{aligned}$$

Using these series expansions in (3.22) and taking the logarithm of that equation, we obtain an expression for the logarithm of the asset return

$$\bar{r} + qr_{1t+1} = \log [\exp(\bar{z} + qz_{1t+1}) + 1] - (\bar{z} + qz_{1t}) + (\bar{g} + qg_{1t+1}). \quad (3.23)$$

We now construct the first-order series expansion of this relationship. The zeroth order term is obtained by evaluating equation (3.23) at $q = 0$:

$$\bar{r} = \log [\exp(\bar{z}) + 1] - \bar{z} + \bar{g},$$

and the relationship for the first-order process is obtained by differentiating (3.23) with respect to q and evaluating the derivative at $q = 0$:

$$r_{1t+1} = \frac{\exp(\bar{z})}{\exp(\bar{z}) + 1} z_{1t+1} - z_{1t} + g_{1t+1}. \quad (3.24)$$

We denote

$$\rho = \frac{\exp(\bar{z})}{\exp(\bar{z}) + 1}.$$

The advantage of the linear relationship (3.24) relative to the original definition (3.21) is that the linear expression can be easily solved forward. Use (3.24) to express the price-dividend z_{1t} , and iterate forward

$$\begin{aligned} z_{1t} &= g_{1t+1} - r_{1t+1} + \rho z_{1t+1} = g_{1t+1} - r_{1t+1} + \rho (g_{1t+2} - r_{1t+2}) + \rho^2 z_{1t+2} = \dots \\ &= \lim_{T \rightarrow \infty} \sum_{j=1}^T \rho^j (g_{1t+j} - r_{1t+j}) + \lim_{T \rightarrow \infty} \rho^T z_{1T}. \end{aligned}$$

Under a transversality (no-bubble) condition

$$\lim_{T \rightarrow \infty} \rho^T z_{1T} = 0,$$

we obtain

$$z_{1t} = \sum_{j=1}^{\infty} \rho^j (g_{1t+j} - r_{1t+j}) = \sum_{j=1}^{\infty} \rho^j g_{1t+j} - \sum_{j=1}^{\infty} \rho^j r_{1t+j}. \quad (3.25)$$

This equation expresses (up to the constant term \bar{z}) the current price-dividend ratio z_{1t} as a function of *future* dividend growth rates g_{1t+j} and *future* returns. This expression is an *accounting identity* derived from the definition of the return (3.24), only using the transversality condition.

The equation states that if the current price-dividend ratio z_{1t} increases, the right-hand side of (3.25) has to increase as well. This increase has to come either in the form of an increase of future dividend growth rates g_{1t+j} , or in the form of a decrease in future returns, or a combination of both.

Since equation (3.25) is an accounting identity, the fact that an increase in the price-dividend ratio has to be accompanied by an increase in the right-hand side of (3.25) does not have any economic content (beyond the ability to test the validity of the transversality condition). In particular, it does not imply that economic causality runs from changes in the current price-dividend ratio z_{1t} to changes in future dividend growth rates and returns.

However, we can study statistical association of price-dividend ratios with future dividend growth and returns to understand which of the two factors on the right-hand side explains more of the variation in price dividend ratio. To do so, compute the covariance of both sides of the equation with z_{1t} . This yields

$$Var(z_{1t}) = Cov\left(z_{1t}, \sum_{j=1}^{\infty} \rho^j g_{1t+j}\right) + Cov\left(z_{1t}, -\sum_{j=1}^{\infty} \rho^j r_{1t+j}\right),$$

and dividing by $Var(z_{1t})$, we obtain

$$1 = \frac{Cov\left(z_{1t}, \sum_{j=1}^{\infty} \rho^j g_{1t+j}\right)}{Var(z_{1t})} + \frac{Cov\left(z_{1t}, -\sum_{j=1}^{\infty} \rho^j r_{1t+j}\right)}{Var(z_{1t})}.$$

Both terms on the right-hand side of this equation can be estimated empirically by truncating the infinite sums. The literature systematically finds that the second term is substan-

tially large than the first term, implying that a majority of movements in the price-dividend ratio are associated with subsequent low returns. An interpretation of this result is that price-dividend ratios do not predominantly move because of fluctuations in expected future dividend growth but rather because of fluctuations in required (or expected) returns, or discount rates. Since expected returns can be decomposed into the risk-free rate and risk premium, and we know that the risk-free rate empirically moves very little, it must be that fluctuations in price-dividend ratios must be largely attributable to fluctuations in risk premia.

Another way how to view this decomposition result is to consider a vector autoregression for the vector

$$x_t = \begin{pmatrix} r_{1t} \\ g_{1t} \\ y_t \end{pmatrix}$$

where y_t is a subvector of variables used to predict returns r_{1t} and dividend growth g_{1t} , perhaps also including lagged variables. Using coordinate vectors, we can express $r_{1t} = e_1' x_t$ and $g_{1t} = e_2' x_t$. Specifying

$$x_{t+1} = A_0 x_t + C w_{t+1}, \quad (3.26)$$

we have

$$\begin{aligned} E_t [r_{1t+j}] &= e_1' E_t [x_{t+j}] = e_1' A_0^j x_t \\ E_t [g_{1t+j}] &= e_2' E_t [x_{t+j}] = e_2' A_0^j x_t \end{aligned}$$

The VAR can be estimated using OLS. Then we can express the expressions on the right-hand side of (3.25) as

$$E_t \left[\sum_{j=1}^{\infty} \rho^j r_{1t+j} \right] = \sum_{j=1}^{\infty} e_1' (\rho A_0)^j x_t = e_1' (\rho A_0) [I - \rho A_0]^{-1} x_t \quad (3.27)$$

$$E_t \left[\sum_{j=1}^{\infty} \rho^j g_{1t+j} \right] = \sum_{j=1}^{\infty} e_2' (\rho A_0)^j x_t = e_2' (\rho A_0) [I - \rho A_0]^{-1} x_t \quad (3.28)$$

At this point, it is useful to review the role of the expectations operator in the above equations. The conditional expectations $E_t [\cdot]$ are taken with respect to the information set generated by the variables in the VAR up to time t , which, given the Markov structure of the VAR, reduces to x_t . However, the full information set at time t is larger if we decide not to include all relevant predictors of returns and dividend growth into the VAR. This means that when we take conditional expectations of equation (3.25) conditional on time- t variables in the VAR, we write

$$E_t [z_{1t}] = E_t \left[\sum_{j=1}^{\infty} \rho^j g_{1t+j} \right] - E_t \left[\sum_{j=1}^{\infty} \rho^j r_{1t+j} \right] = (e_2 - e_1)' (\rho A_0) [I - \rho A_0]^{-1} x_t.$$

The left-hand side expression $E_t [z_{1t}]$ is a projection of the price-dividend ratio z_{1t} on the time- t VAR variables. It can be equal to z_{1t} if z_{1t} itself is perfectly predicted by x_t . In order to understand how the contribution of a predictor variable that moves the price-dividend ratio decomposes into movements in future returns and future dividend growth, we can separately evaluate expressions (3.27)–(3.28).

Empirically, future dividend growth is much harder to predict than future returns. To illustrate the consequence of this point forcefully, consider a simple example in which dividend growth is iid, while returns can be predicted using a persistent scalar variable y_t :

$$\begin{aligned} r_{1t+1} &= y_t + \sigma_r w_{t+1}^r \\ g_{1t+1} &= \sigma_g w_{t+1}^g \\ y_{t+1} &= \phi y_t + \sigma_y w_{t+1}^y \end{aligned}$$

where w_{t+1}^r , w_{t+1}^g and w_{t+1}^y are independent shocks. This problem can be mapped into the VAR (3.26) with matrices

$$A_o = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & \phi \end{pmatrix} \quad C = \begin{pmatrix} \sigma_r & 0 & 0 \\ 0 & \sigma_g & 0 \\ 0 & 0 & \sigma_y \end{pmatrix}$$

with $x_t = (r_{1t}, g_{1t}, y_t)'$ and $w_{t+1} = (w_{t+1}^r, w_{t+1}^g, w_{t+1}^y)'$. Solving for expressions (3.27)–(3.28) yields

$$E_t \left[\sum_{j=1}^{\infty} \rho^j r_{1t+j} \right] = \frac{\rho}{1 - \rho\phi} y_t \quad E_t \left[\sum_{j=1}^{\infty} \rho^j g_{1t+j} \right] = 0$$

and hence

$$z_{1t} = -\frac{\rho}{1 - \rho\phi} y_t.$$

Changes in expected returns thus drive the price-dividend ratios, and the sensitivity of the price-dividend ratio to the predictor variable is the larger the more persistent the predictor is.

3.2.4 Linear approximation of valuation ratios

We now study a log-linear approximation of the price dividend ratio in a valuation model. Consider the valuation equation

$$\frac{Q_t}{G_t} = E_t \left[\frac{S_{t+1}}{S_t} \frac{G_{t+1}}{G_t} \left(\frac{Q_{t+1}}{G_{t+1}} + 1 \right) \right].$$

Contrary to the accounting identity in the form of definition of returns (3.21) that we worked with in the previous subsection, this valuation equation, which links price-dividend ratios, dividend growth and the stochastic discount factor, is a model that predicts the price-dividend ratio given a model of the SDF and cash flows.

Denote

$$z_t = \log \frac{Q_t}{G_t} \quad s_{t+1} = \log \frac{S_{t+1}}{S_t} \quad g_{t+1} = \log \frac{G_{t+1}}{G_t},$$

which allows to rewrite the valuation equation as.

$$\exp(z_t) = E_t [\exp(s_{t+1} + g_{t+1}) (\exp(z_{t+1}) + 1)]$$

Now assume that each of the processes q_t, s_t, g_t can be written in the series expansion form (for perturbation parameter q)

$$\exp(\bar{z} + qz_{1t}) = E_t [\exp(\bar{s} + qs_{1t+1} + \bar{g} + qg_{1t+1}) (\exp(\bar{z} + qz_{1t+1}) + 1)].$$

Computing the zeroth and first-order derivative of this equation with respect to q yields

$$\begin{aligned} \exp(\bar{z}) &= \exp(\bar{s} + \bar{g}) (\exp(\bar{z}) + 1) \\ \exp(\bar{z}) z_{1t} &= E_t [\exp(\bar{s} + \bar{g}) (s_{1t+1} + g_{1t+1}) (\exp(\bar{z}) + 1) + \exp(\bar{s} + \bar{g}) \exp(\bar{z}) z_{1t+1}]. \end{aligned}$$

Using the steady-state equation from the first line in the equation for the first-derivative process on the second line, we get

$$z_{1t} = E_t [s_{1t+1} + g_{1t+1} + \exp(\bar{s} + \bar{g}) z_{1t+1}]. \quad (3.29)$$

This is a linear forward-looking equation for the price-dividend ratio, or, more specifically, for the deviations of the logarithm of the price-dividend ratio from the steady state value. In order to solve this equation, let us impose linear dynamics on the model. Specifically, assume a linear law of motion for the state $x_t \in \mathbb{R}^n$

$$x_{t+1} = A_0 x_t + C w_{t+1}, \quad w_{t+1} \sim N(0, I_p)$$

and impose a linear structure of the SDF and dividend growth rate

$$\begin{aligned} s_{1t+1} &= S x_{t+1} \\ g_{1t+1} &= G x_{t+1} \end{aligned}$$

where S and G are $1 \times n$ vectors.

With this specification, we can conjecture that the solution for the price-dividend ratio is also linear

$$z_{1t} = Q x_t$$

with Q being a $1 \times n$ coefficient vector to be solved for. Plugging these specifications into the valuation equation yields

$$\begin{aligned} Q x_t &= E_t [S x_{t+1} + G x_{t+1} + \exp(\bar{s} + \bar{g}) Q x_{t+1}] \\ &= (S + G) A_0 x_t + \exp(\bar{s} + \bar{g}) Q A_0 x_t \end{aligned}$$

This equation has to hold for every value of x_t , so coefficients on the left- and right-hand

side have to match

$$Q = (S + G) A_o + \exp(\bar{s} + \bar{g}) Q A_o.$$

We can therefore solve for the vector Q

$$Q = (S + G) A_o [I - \exp(\bar{s} + \bar{g}) A_o]^{-1}$$

which represents the present discounted value of future dividends under the log-linearized model.

A key question is what the log-linear approximation of the valuation equation omits. A central observation is that under this approximation, the mapping between the price-dividend ratio z_{1t} and the state of the economy x_t

$$z_{1t} = Qx_t = (S + G) A_o [I - \exp(\bar{s} + \bar{g}) A_o]^{-1} x_t$$

does not depend on uncertainty of the underlying state, embedded in the parameter matrix C . Increasing the volatility of shock by increasing the magnitude of C will therefore make the state, and hence the price-dividend ratio, more volatile but will not generate any compensation for risk in the form of higher risk premia that would push the average price-dividend ratio lower.

To understand why this is the case, recall that risk premia are given by covariances of the stochastic discount factor with returns, since the Euler equation

$$0 = E_t \left[\frac{S_{t+1}}{S_t} (R_{t+1} - R_{t+1}^f) \right]$$

implies

$$E_t [R_{t+1} - R_{t+1}^f] = -R_{t+1}^f \text{Cov}_t \left[\frac{S_{t+1}}{S_t}, R_{t+1} - R_{t+1}^f \right].$$

The linear approximation of the valuation of

$$z_{1t} = E_t [s_{1t+1} + g_{1t+1} + \exp(\bar{s} + \bar{g}) z_{1t+1}].$$

neglects these covariances because of the additive form between the SDF and future dividends and price-dividend ratios.

This conclusion is immediately related to the fact that compensation for risk in a model of smooth preferences with risk aversion characterized by the [Arrow \(1965\)](#)–[Pratt \(1965\)](#) measure is a ‘second-order’ concept. To see this result more clearly, let us revisit the series expansion of the valuation equation

$$\exp(\bar{z} + qz_{1t}) = E_t [\exp(\bar{s} + qs_{1t+1} + \bar{g} + qg_{1t+1}) (\exp(\bar{z} + qz_{1t+1}) + 1)].$$

But before we log-linearize the equation, let us first manipulate the expression on the right-hand side. Substitute in the linear model

$$s_{1t+1} = Sx_{t+1} \quad g_{1t+1} = Gx_{t+1} \quad x_{t+1} = A_o x_t + Cw_{t+1}$$

together with the linear conjecture for the price dividend ratio $z_{1t+1} = Qx_{t+1}$ to obtain

$$\begin{aligned} \exp(\bar{z} + qz_{1t}) &= E_t [\exp(\bar{s} + \bar{g} + \bar{q} + q(S + G + Q)(A_o x_t + Cw_{t+1}))] \\ &\quad + E_t [\exp(\bar{s} + \bar{g} + q(S + G)(A_o x_t + Cw_{t+1}))]. \end{aligned}$$

Collecting the deterministic and random components, we obtain

$$\begin{aligned} \exp(\bar{z} + qz_{1t}) &= \exp(\bar{s} + \bar{g} + \bar{q} + q(S + G + Q)A_o x_t) E_t [\exp(q(S + G + Q)Cw_{t+1})] \\ &\quad + \exp(\bar{s} + \bar{g} + q(S + G)A_o x_t) E_t [\exp(q(S + G)Cw_{t+1})]. \end{aligned}$$

In order to compute the expectation of the exponential, we utilize the formula for the expectation of a log-normally distributed random variable

$$w \sim N(0, I) \quad \implies \quad E[\exp(\mu + \sigma w)] = \exp\left(\mu + \frac{1}{2}\sigma\sigma'\right).$$

In our case, $\sigma = q(S + G + Q)$ and $\sigma = q(S + G)$, and hence

$$\begin{aligned} \exp(\bar{z} + qz_{1t}) &= \exp(\bar{s} + \bar{g} + \bar{q} + q(S + G + Q)A_o x_t) \cdot \\ &\quad \cdot \exp\left(\frac{1}{2}q^2(S + G + Q)CC'(S + G + Q)'\right) \\ &\quad + \exp\left(\bar{s} + \bar{g} + q(S + G)A_o x_t + \frac{1}{2}q^2(S + G)CC'(S + G)'\right) \end{aligned}$$

The effect of risk premia is embedded in in the two terms involving the covariance CC' . These two terms interact the uncertainty in the SDF, cash flows, and next-period price-dividend ratio, and scale linearly with the covariance of the shocks CC' , rather than linearly with volatility.

More importantly from the perspective of the perturbation approximation, the risk premium contribution scales with q^2 , so it vanishes in the linear approximation. We need at least a second-order approximation of the valuation equation to pick up the risk premium effect, or rely on an alternative perturbation technique. For example, [Borovička and Hansen \(2014\)](#), and [Bhandari et al. \(2019\)](#) utilize a series expansion in which the perturbation parameter q scales down the volatility of the shocks while at the same time scaling up the risk aversion parameter in the preference specification, so that compensation for risk in the form of risk premia does not vanish in the linear approximation.

3.2.5 Linear-quadratic approximation of stochastic control models

We now focus on the approximation of a constrained stochastic control model. The model can be written in sequence form as

$$\max_{\{a_t\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t u(x_t, a_t) \quad \text{subject to } x_{t+1} = \psi(x_t, a_t, w_{t+1})$$

where x_t is the state vector, with x_0 given, a_t is the vector of controls that affect utility and the evolution of the state, and the function ψ represents the set of restrictions that determine the controlled law of motion for the state.

We again apply the series expansion method. The approximation we use produces a first-order approximation of the controlled law of motion, and a second-order approximation of the objective function. With this approximation, we turn the problem into a linear-quadratic control problem that we studied in Section 3.1.3, and that offers semi-analytical solutions in the form of a linear policy function.

We start with the first-order series expansion of the controlled law of motion. We approximate the state and control dynamics to first order as

$$x_t(\mathbf{q}) \approx \bar{x} + \mathbf{q}x_{1t} \quad a_t(\mathbf{q}) \approx \bar{a} + \mathbf{q}a_{1t}$$

and write the class of laws of motion indexed by the perturbation parameter as

$$x_{t+1}(\mathbf{q}) = \psi(x_t(\mathbf{q}), a_t(\mathbf{q}), \mathbf{q}w_{t+1}, \mathbf{q})$$

The first-order approximation of the controlled law of motion then is

$$\begin{aligned} \bar{x} &= \psi(\bar{x}, \bar{a}, 0) \\ x_{1t+1} &= \psi_q + \psi_x x_{1t} + \psi_a a_{1t} + \psi_w w_{t+1} \end{aligned}$$

The partial derivative matrices ψ_q , ψ_x , ψ_a , and ψ_w are evaluated at the steady state $(\bar{x}, \bar{a}, 0, 0)$.

With the linear approximation of the state and control, we approximate the period utility function $u(x_t, a_t)$ to second order. We write the period utility process as

$$u(x_t(\mathbf{q}), a_t(\mathbf{q}), \mathbf{q}) = u_t(\mathbf{q}) \approx \bar{u} + \mathbf{q}u_{1t} + \frac{\mathbf{q}^2}{2}u_{2t} \quad (3.30)$$

In order to obtain the derivative processes \bar{u} , u_{1t} , and u_{2t} , evaluate the derivatives with respect to \mathbf{q} of

$$u(x_t(\mathbf{q}), a_t(\mathbf{q}), \mathbf{q}) \approx u(\bar{x} + \mathbf{q}x_{1t}, \bar{a} + \mathbf{q}a_{1t}, \mathbf{q}).$$

The zeroth-order derivative evaluates the period utility function at the steady state

$$\bar{u} = u(\bar{x}, \bar{a}, 0).$$

The first-order derivative of the utility function is given by

$$u_{1t} = u_q + u_x x_{1t} + u_a a_{1t}.$$

Finally, the second-order derivative is

$$u_{2t} = u_{qq} + 2u_{qx}x_{1t} + 2u_{qa}a_{1t} + x'_{1t}u_{xx}x_{1t} + a'_{1t}u_{aa}a_{1t} + 2a'_{1t}u_{ax}x_{1t}$$

where all partial derivatives of u are evaluated at the steady state $(\bar{x}, \bar{a}, 0)$

Then we can construct the second-order approximation of u_t in (3.30) evaluated at $q = 1$, by combining \bar{u} , u_{1t} and $\frac{1}{2}u_{2t}$:

$$u(x_t, a_t) \approx \bar{u} + u_{1t} + \frac{1}{2}u_{2t}.$$

We thus obtain the decision problem

$$\max_{\{a_{1t}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t \hat{u}(x_t, a_t)$$

with

$$\begin{aligned} \hat{u}(x_t, a_t) = & \bar{u} + u_q + \frac{1}{2}u_{qq} + (u_x + u_{qx})x_{1t} + (u_a + u_{qa})a_{1t} \\ & + \frac{1}{2}x'_{1t}u_{xx}x_{1t} + \frac{1}{2}a'_{1t}u_{aa}a_{1t} + a'_{1t}u_{ax}x_{1t} \end{aligned}$$

subject to

$$x_{1t+1} = \psi_q + \psi_x x_{1t} + \psi_a a_{1t} + \psi_w w_{t+1}$$

with a given initial condition x_0 . This is a linear-quadratic problem with a tractable solution even for high-dimensional state spaces. The solution of the problem is provided in Section 3.1.3. The optimal control a_t^* is a linear function of the state x_t , the solution is based on computation of quadratic sums using formula (3.11) together with an optimization step.

3.2.6 Solving linear equilibrium models

Dynamic equilibria in macroeconomic and asset pricing modeling often feature a combination of backward-looking and forward-looking equations.

A backward-looking equation is characterized by an initial state, and represents the current value of the variable as a function of past realizations or innovations, for example in the form of a law of motion

$$x_{t+1} = A_0 x_t + C w_{t+1}.$$

On the contrary, a forward-looking equation represents the current value of a variable as a function of its future values, for example in the form of an Euler equation

$$z_{1t} = E_t [s_{1t+1} + g_{1t+1} + \exp(\bar{s} + \bar{g}) z_{1t+1}].$$

The path of a forward-looking equation is pinned down by a transversality condition.

We want to find a solution of the dynamic model in the form of a law of motion for a set of variables that define the appropriate notion of the state vector, which will be backward-looking, with an appropriate initial condition, and a mapping from the state to all remaining endogenous variables.

As an example, let us consider the valuation model for the price-dividend ratio. In that

model, we have the backward-looking law of motion of the state vector that pins down the model for the SDF and dividend growth

$$x_{t+1} = A_o x_t + C w_{t+1} \quad s_{1t+1} = S x_{t+1} \quad g_{1t+1} = G x_{t+1},$$

and the forward-looking equation

$$z_{1t} = E_t [s_{1t+1} + g_{1t+1} + \exp(\bar{s} + \bar{g}) z_{1t+1}]$$

which we solved forward, imposing the appropriate transversality condition

$$\lim_{T \rightarrow \infty} \exp(T(\bar{s} + \bar{g})) z_{1t+T} = 0.$$

to obtain the price-dividend ratio as a function of the state, $z_{1t} = Q x_t$.

Another example is the deterministic neoclassical growth model. The backward-looking equation is the law of motion for capital, controlled by the consumption-investment choice

$$k_{t+1} = (1 - \delta) k_t + G(k_t) - c_t,$$

with a given initial condition k_0 . The forward-looking equation is the Euler equation for the optimal consumption choice.

$$U'(c_t) = \beta U'(c_{t+1}) (1 - \delta + G'(k_{t+1}))$$

This equation needs to be solved forward, and a transversality condition imposed that states that the optimal choice of consumption leads the state variable to the steady state. The solution of the problem involves finding the mapping $c_t = c(k_t)$.

There are well-established methods for solving these sets of linear forward- and backward-looking equations, and the conditions stated in [Blanchard and Kahn \(1980\)](#), and [Sims \(2002\)](#) spell out required conditions for a unique stable solution to exist. [Ljungqvist and Sargent \(2018\)](#), Chapter 5 provide a pedagogical treatment. Implementations of these model solutions are available in existing software packages, in particular in Dynare.

3.2.7 VAR estimation and identification of structural shocks

Another appealing feature of the linear vector autoregression setup is tractability of estimation. Imagine we have a law of motion for the a vector of macroeconomic variables in the form of a VAR

$$x_{t+1} = A_o x_t + C w_{t+1}$$

where x_t are the variables of interest, and w_{t+1} consists of a vector so-called 'structural' macroeconomic shocks, i.e., shocks that have a particular interpretation in the context of macroeconomic models, like shocks to total factor productivity or shocks emerging from monetary or fiscal policy.

The matrix A_o can be estimated using OLS, equation by equation, with each equation

estimating one row of the matrix. In order to estimate the covariance matrix of the innovations, $\Sigma = CC'$, we then collect residuals from the estimation of individual rows of the VAR and compute their empirical covariance.

A complication arises when we want to obtain C from the knowledge of $\Sigma = CC'$. There are generally many solutions to this decomposition, and we need to impose additional restrictions to be able to identify C . These additional restrictions must come from the economic structure of the problem. This problem is known as shock identification problem in structural VARs.

To understand the problem, imagine we want to understand the dynamic impact of a shock on the macroeconomy. Recall the impulse response function $[h_t]_{ik} = [A_o^t C]_{ik}$ from Section 3.1.2 that measures the impact of the k -th component of the time-0 shock vector, w_0^k , on the i -th component of the vector of time- t macroeconomic variables, x_t^i . We need to know the matrix C to be able to infer the impulse response.

Now imagine that instead of decomposing Σ as CC' , we find another decomposition $\Sigma = \tilde{C}\tilde{C}'$. With the alternative decomposition, we would now incorrectly measure the impulse response as $[A_o^t \tilde{C}]_{ik}$, so instead of isolating the true response of the particular shock w_0^k , we would infer some linear combination of impulse responses of all the shocks.

One of the many examples of imposing structure that helps with identification is the so-called Cholesky decomposition, which is useful if we have a specific idea on the lags with which particular shocks affect the variables of interest. The Cholesky decomposition produces a lower triangular matrix C .

Christiano et al. (2005) use this structure to identify the effects of monetary policy surprises on macroeconomic variables. Imagine that x_t consists of inflation π_t , real GDP y_t and nominal interest rate i_t . We want to study the effect of an increase in i_t on inflation and GDP. The obvious concern is endogeneity—the Federal Reserve adjusts i_t in response to changes in macroeconomic conditions. Christiano et al. (2005) therefore make the identification assumption that states that an exogenous change in the interest rate is such that inflation and output cannot respond to it contemporaneously in the same period that the shock to the interest rate occurs.

This identification assumption can be implemented as follows. We order i_t last in the VAR vector x_t , and use a Cholesky decomposition of CC' . With this structure, we have

$$\begin{pmatrix} \pi_{t+1} \\ y_{t+1} \\ i_{t+1} \end{pmatrix} = A_o(\theta) \begin{pmatrix} \pi_t \\ y_t \\ i_t \end{pmatrix} + \begin{pmatrix} C_{11} & 0 & 0 \\ C_{21} & C_{22} & 0 \\ C_{31} & C_{32} & C_{33} \end{pmatrix} \begin{pmatrix} w_{t+1}^1 \\ w_{t+1}^2 \\ w_{t+1}^3 \end{pmatrix}.$$

The last component of the shock vector, w_{t+1}^3 , is the identified exogenous monetary policy shock. When it occurs in period 0, it affects i_0 but not π_0 or y_0 . Future values of π_t and y_t are only affected through the propagation encoded in the matrix A_o , and this impact is captured in the impulse response function $[h_t]_{i3} = [A_o^t C]_{i3}$ for $i = 1, 2$.

Another frequent situation is the case when A_o and C are matrices that are not unconstrained but have a particular structure implied by an economic model. Specifically,

consider a vector of structural parameters θ that restricts the entries of the matrices A_o and C , such that A_o and C are functions of θ :

$$x_{t+1} = A_o(\theta) x_t + C(\theta) w_{t+1}.$$

We could still estimate the $A_o(\theta)$ in an unconstrained way consistently using OLS but if we want to explicitly impose the cross-equation restrictions embedded in $A_o(\theta)$ and $C(\theta)$ by estimating θ directly, we cannot rely on OLS anymore. A variety of estimating techniques can then be used instead, including GMM, maximum likelihood methods, or Bayesian estimation.

3.2.8 A cautious note on the accuracy of perturbation approximations

Since the solution to the linear-quadratic problem in Section 3.2.5 or linearized equilibrium models in Section 3.2.6 only involves operations with coefficient matrices, it is highly scalable at negligible computational costs. However, by construction, it provides an approximation that is designed to work well if the economy resides in the neighborhood of the steady state and the nonlinearities are not substantial.

Importantly, there is a critical distinction in terms of judging the accuracy of the approximation relative to the intuition provided by Taylor's theorem discussed in Section 3.2.1. The Taylor's theorem implies that accuracy of the linear or polynomial approximation of a smooth nonlinear function in the neighborhood of the expansion only depends on the local properties (local curvature) of that function, not on the behavior of the function far from the expansion point.

Because of forward-looking behavior in dynamic stochastic models, this intuition does not carry over to these models. Even in the neighborhood of the steady state, the endogenous decisions depend on expectations of future realizations across the whole state space. So when the approximation substantially alters the behavior of the model in the tails of the state space, this will have also have consequences for the behavior near the steady state. In this sense, the accuracy of the approximation is no longer 'localized'. Evaluation of this accuracy therefore becomes much more involved and requires a good understanding of the extent to which these inaccuracies due to anticipation effects are sufficiently modest.

Chapter 4

Filtering and learning

Textbook: Ljungqvist and Sargent (2018), Chapter 2 (Sections 2.7–2.9, Kalman filter).
Applications: Jovanovic (1979), Muth (1960).
QuantEcon: Quantitative Economics with Python, Topic 31 (Kalman filter), Topic 50–52 (Bayes law), Topic 53 (search with learning)

In this chapter, we study a problem where the economic agent or econometrician does not have a perfect information about the state of the economy and must learn about it from noisy observations. We follow the seminal work by Kalman (1960) and study a problem where the underlying state follows a vector autoregression with Gaussian innovations, and the measurement is a linear function of the state with added measurement noise. With a Gaussian prior and Gaussian shocks, the posterior will also be Gaussian. In the language of Bayesian probability theory, a Gaussian prior is a conjugate prior for a likelihood function of that data that is also Gaussian.

Under this conjugate prior structure, the inference problem becomes one in which we only need to deduce expressions for the mean and variance of the posterior because a Gaussian distribution is uniquely pinned down by these two moments. Kalman (1960) derived a recursive formula for these expressions.

4.1 Kalman filter

We now assume that the realizations of linear vector autoregression are not observable, and the econometrician observes instead a noisy measurement of the current state. The problem, solved by Kalman (1960), is to construct the optimal forecast of the path of the underlying state, given the observed measurements. The system can be written as

$$\begin{aligned} x_{t+1} &= A_0 x_t + C w_{t+1} & n \times 1 & & w_{t+1} &\sim N(0, I_p) \\ y_t &= G x_t + v_t & m \times 1 & & v_t &\sim N(0, R). \end{aligned} \tag{4.1}$$

Here, x_t is an unobservable state vector that follows a Gaussian Markov process, y_t is a measurement vector and v_t is measurement noise independent of $\{w_t\}_{t=1}^{\infty}$. We assume an initial condition

$$x_0 \sim N(\hat{x}_0, \Sigma_0). \quad (4.2)$$

The algorithm leads to a recursive formula for the best predictor of x_t given observations $y^{t-1} = (y_{t-1}, \dots, y_0)$. Because of the linear-Gaussian structure of the whole system, we can infer that the predictor will also be Gaussian, so it is sufficient to construct predictors for the first two moments of its distribution

$$\begin{aligned} \hat{x}_t &= E[x_t | y^{t-1}] \\ \Sigma_t &= E[(x_t - \hat{x}_t)(x_t - \hat{x}_t)']. \end{aligned} \quad (4.3)$$

Since the underlying state x_t is persistent, observed data y_t will be serially correlated. The idea is to derive the contribution of the information embedded in the new observation y_t to the prediction of x_{t+1} , relative to what we already could have inferred from y^{t-1} . Mathematically, we are applying the **Gram-Schmidt orthogonalization process** to the sequence of the data observations y^t , and constructing an orthogonal basis of **innovations** (a_0, a_1, \dots, a_t) from (y_0, y_1, \dots, y_t) . A particular innovation a_t will represent new information in the additional data point y_t relative to what we learned from y^{t-1} . Since the innovations are orthogonal, projecting the unobserved state x_t on the innovations will be simpler than projecting on the original data. Effectively, we are replacing (4.3) with an equivalent representation

$$\hat{x}_t = E[x_t | a^{t-1}].$$

Observe that if we are successful, the method will constitute a massive simplification of the inference problem in this class of models. We have seen in Section 1.1.7 that filtering the Markov state on a hidden Markov chain involved a recursive formula for the whole distribution of the state. Here, the recursive formula reduces to keeping track of the first two moments.

It turns out that the Kalman filtering problem and the problem of iteratively solving a Bellman equation in a linear-quadratic environment, studied in Section 3.1.3 are related through a reversal of the time axis. While the Kalman filter proceeds forward in time, the dynamic programming method on the Bellman equation proceeds backwards.

Let us start from the initial condition (4.2) and study how much we can learn about the unknown x_0 from observing y_0 . We will proceed by using linear projections, which corresponds to running theoretical OLS regressions. This method is justified as a way of obtaining efficient forecasts given the linear-Gaussian environment for the problem.

Starting from the initial distribution for x_0 , we infer

$$y_0 \sim N(G\hat{x}_0, G\Sigma_0G' + R).$$

To construct the forecast of x_0 given the data point y_0 , let us project the unknown $x_0 - \hat{x}_0$

on the new information embedded in the observation of y_0 :

$$x_0 - \hat{x}_0 = L_0 \underbrace{(y_0 - G\hat{x}_0)}_{\text{innovation } a_0} + \eta_0$$

The residual η_0 is orthogonal on the innovation $a_0 \doteq y_0 - G\hat{x}_0$ by construction. The innovation a_0 represents the ‘surprise’, or new information, embedded in the observation of y_0 relative to its expected value $G\hat{x}_0$. This innovation also constitutes the first element in the construction of the orthogonal basis (a_0, a_1, \dots) constructed from observations (y_0, y_1, \dots) .

Given the orthogonality between η_0 and a_0 , post-multiplying by $(y_0 - G\hat{x}_0)'$ and taking expectations yields

$$\begin{aligned} E[(x_0 - \hat{x}_0)(y_0 - G\hat{x}_0)'] &= L_0 E[(y_0 - G\hat{x}_0)(y_0 - G\hat{x}_0)'] \\ \Sigma_0 G' &= L_0 (G\Sigma_0 G' + R), \end{aligned}$$

and hence the $n \times m$ matrix regression coefficient L_0 takes the form

$$L_0 = \Sigma_0 G' (G\Sigma_0 G' + R)^{-1}.$$

Equation (4.1) then implies that we can write

$$x_1 = A_o x_0 + Cw_1 = A_o \hat{x}_0 + A_o (x_0 - \hat{x}_0) + Cw_1. \quad (4.4)$$

The mean forecast of the state x_1 given the data point y_0 is therefore given by

$$\begin{aligned} \hat{x}_1 &= E[x_1 | y^0] = E[x_1 | a^0] \\ &= A_o \hat{x}_0 + A_o L_0 (y_0 - G\hat{x}_0) \\ &= A_o \hat{x}_0 + K_0 (y_0 - G\hat{x}_0), \end{aligned} \quad (4.5)$$

where the matrix

$$K_0 = A_o \Sigma_0 G' (G\Sigma_0 G' + R)^{-1}$$

is called the **Kalman gain**. Subtracting (4.5) from (4.4) yields

$$\begin{aligned} x_1 - \hat{x}_1 &= A_o (x_0 - \hat{x}_0) + Cw_1 - K_0 (y_0 - G\hat{x}_0) \\ &= (A_o - K_0 G) (x_0 - \hat{x}_0) + Cw_1 - K_0 v_0. \end{aligned}$$

Notice that the three terms on the previous line are independent. Hence the variance of the forecast given the data point y^0

$$\begin{aligned} \Sigma_1 &= E[(x_1 - \hat{x}_1)(x_1 - \hat{x}_1)'] = \\ &= (A_o - K_0 G) \Sigma_0 (A_o - K_0 G)' + CC' + K_0 R K_0' \end{aligned}$$

We therefore have the distribution $x_1 | y^0 \sim N(\hat{x}_1, \Sigma_1)$. We therefore have the recursive

system

$$\begin{aligned}
 a_t &= y_t - G\hat{x}_t & (4.6) \\
 K_t &= A_0\Sigma_t G' (G\Sigma_t G' + R)^{-1} \\
 \hat{x}_{t+1} &= A_0\hat{x}_t + K_t a_t \\
 \Sigma_{t+1} &= (A_0 - K_t G)\Sigma_t (A_0 - K_t G)' + CC' + K_t R K_t'
 \end{aligned}$$

- The first equation defines the **innovation** a_t , which is the deviation of the observed y_t from its best predictor $G\hat{x}_t$ constructed given y^{t-1} .
- The second equation defines the **Kalman gain**, which tells how much the innovation updates the previous best guess of the state \hat{x}_{t+1} .
- The third equation is the law of motion for the mean forecast \hat{x}_{t+1} . Notice that the best forecast of \hat{x}_{t+1} given y^{t-1} is $A_0\hat{x}_t$, to which we add $K_t a_t$ as the contribution of the information from y_t .
- Finally, we update the accuracy (variance) of the forecast Σ_{t+1} .

We can substitute for K_t into the law of motion for Σ_t to obtain

$$\Sigma_{t+1} = A_0\Sigma_t A_0' - A_0\Sigma_t G' (G\Sigma_t G' + R)^{-1} G\Sigma_t A_0' + CC'. \quad (4.7)$$

This is a matrix **Riccati equation** which often appears in linear-quadratic dynamic programming.

- Observe that while the evolution of \hat{x}_t is stochastic, being updated by the innovations that are constructed from the observations of y_t , the evolution of Σ_t is deterministic. Σ_t typically converges to a constant in a time-invariant model, and the constant is zero when $CC' = 0$. This says that all observations y_t are equally informative, regardless of their particular value. This result is specific to this particular model.
- The path \hat{x}_t is often called the **filtered path** of x_t , in other words it represents the most likely location of x_t conditional on y^{t-1} .
- Can subsequent realizations of y_{t+j} , $j = 0, 1, 2, \dots$ make the estimate of x_t more precise? They can. y_{t+j} is a signal about x_{t+j} , and knowledge where the state x is at time $t + j$ is also informative about where the state has been at time t . This is what the **Kalman smoother** does.

4.2 Applications

4.2.1 Muth's example

In the 1950's, Phillip Cagan (Cagan (1956)), Milton Friedman (Friedman (1957)), and others studied models of **adaptive expectations**, in which expectations about the future slowly

adjust in response to arrival of new data. In Muth (1960), John Muth asked what type of underlying stochastic processes would ‘rationalize’ the adaptive expectations model as the best statistical forecast of the future. In this way, he wanted to identify agent’s beliefs (forecasts) with objective (mathematical) expectations, given available information. This can be viewed a precursor of the assumption of rational expectations, more fully developed in Muth (1961).

The solution to the problem is close to the filtering solution of Kalman (1960), despite the fact that Muth did not use the Kalman filter directly. The idea is to postulate a stochastic process under which the adaptive expectations model can be interpreted as the result of optimal learning (filtering).

Specifically, consider the model for agent’s adaptive expectations

$$\begin{aligned} y_{t+1}^* &= K \sum_{j=0}^{\infty} (1-K)^j y_{t-j} \\ &= (1-K)y_t^* + Ky_t \end{aligned} \quad (4.8)$$

where K is the weight on the current observation for the time- t forecast of y_{t+1} , denoted y_{t+1}^* . Cagan (1956) used this model as a model of agent’s forecasts of future inflation, Friedman (1957) for forecasts of future income.

Muth (1960) studied a model that can be written as a special case of the system in the Kalman filter problem:

$$\begin{aligned} x_{t+1} &= x_t + w_{t+1} \\ y_t &= x_t + v_t \end{aligned} \quad (4.9)$$

where w_t, v_t are independent scalar shocks with covariances Q and R , respectively, and y_t and x_t are also scalar. In the context of the Kalman filter model (4.1), we have

$$A_0 = 1, CC' = Q, G = 1$$

Then the filtering equations (4.6) together with (4.7) become

$$\begin{aligned} a_t &= y_t - \hat{x}_t \\ K_t &= \frac{\Sigma_t}{\Sigma_t + R} \\ \hat{x}_{t+1} &= \hat{x}_t + K_t a_t \\ \Sigma_{t+1} &= \Sigma_t - \frac{\Sigma_t^2}{\Sigma_t + R} + Q \end{aligned}$$

When we take the limit as $t \rightarrow \infty$, we expect $\Sigma_t \rightarrow \Sigma$ and $K_t \rightarrow K$. Then the law of motion

for the forecast is given by

$$\begin{aligned}\hat{x}_{t+1} &= \hat{x}_t + Ka_t = \hat{x}_t + K(y_t - \hat{x}_t) = \\ &= (1 - K)\hat{x}_t + Ky_t \\ &= \frac{R}{\Sigma + R}\hat{x}_t + \frac{\Sigma}{\Sigma + R}y_t\end{aligned}$$

Observe that this forecasting formula is in line with the ‘adaptive’ forecast model (4.8). We can then conclude that optimal filtering (best forecast) in the model (4.9) yields an belief updating formula which Cagan (1956) and Friedman (1957) interpreted as adaptive expectations.

4.2.2 Jovanovic’s model

Jovanovic (1979) has a model of learning about the quality of an employment match between a worker and a firm. The true underlying quality is θ , and the worker-firm pair observes a sequence of output realizations which are noisy realizations of the match quality:

$$\begin{aligned}\theta_{t+1} &= \theta_t = \theta \\ y_t &= \theta_t + v_t\end{aligned}$$

The pair has a prior $\theta \sim N(m_{-1}, \Sigma_0)$. The time- t forecast of θ is denoted $m_t = \hat{x}_{t+1} = E[\theta | y^t]$. The model then fits into the Kalman filter framework with $A_0 = 1$, $C = 0$, $G = 1$, $R > 0$, and we thus obtain

$$\begin{aligned}a_t &= y_t - m_{t-1} \\ K_t &= \frac{\Sigma_t}{\Sigma_t + R} \\ m_t &= m_{t-1} + K_t a_t \\ \Sigma_{t+1} &= \frac{\Sigma_t R}{\Sigma_t + R}.\end{aligned}$$

This can be summarized as

$$\begin{aligned}m_t &= (1 - K_t)m_{t-1} + K_t y_t \\ K_t &= \frac{\Sigma_t}{\Sigma_t + R} \\ \frac{1}{\Sigma_{t+1}} &= \frac{1}{R} + \frac{1}{\Sigma_t}\end{aligned}$$

The quantity Σ_t^{-1} is called precision, and as $\Sigma_t^{-1} \rightarrow \infty$ over time, the value of the parameter θ is ultimately learned. Also, over time, the Kalman gain declines to zero, as additional observations become less and less informative. This is contrary to the case when x_t fluctuates over time.

4.2.3 Innovations representation and VARs

Let us return to the state-space model

$$\begin{aligned}x_{t+1} &= A_0 x_t + C w_{t+1} \\ y_t &= G x_t + v_t\end{aligned}\tag{4.10}$$

We can think about this system as a representation of a Markov equilibrium of a particular economic model. The $p + m$ shocks w_t and v_t have a so-called ‘structural interpretation’ and represent shocks to technology, preferences, endowments, information sets, measurements, etc. However, the Markov state consists of variables x_t which may not have directly observable counterparts.

Instead, the measurement equation provides partial information about x_t . The **innovations representation**

$$\begin{aligned}\hat{x}_{t+1} &= A_0 \hat{x}_t + K_t a_t \\ y_t &= G \hat{x}_t + a_t\end{aligned}$$

has a very different interpretation than the structural system (4.10). Uncertainty is now summarized by m shocks (innovations) a_t which may not have a direct ‘structural’ counterpart. Even under the filtration of the econometrician which is constructed from histories of observable data y^t , these innovations are cross-correlated:

$$\text{Var} \left[a_t | y^{t-1} \right] = \text{Var} \left[G(x_t - \hat{x}_t) + v_t | y^{t-1} \right] = G \Sigma G' + R.$$

Further, while these innovations are independent over time under the filtration of the econometrician, they are dependent under the complete information generated by the data-generating process and combine structural shocks w_t and measurement errors v_t :

$$\begin{aligned}a_t &= y_t - G \hat{x}_t = G(x_t - \hat{x}_t) + v_t = \\ &= GA(0x_{t-1} - \hat{x}_{t-1}) - GK_{t-1}a_{t-1} + GCw_t + v_t.\end{aligned}$$

Nevertheless, the innovations a_t can be more informative about structural shocks w_t if we can impose more structure on the problem. We are therefore interested in methods that could help us extract these innovations without having to estimate the whole state-space representation of the model.

It turns out that even though the observations y_t do not have a Markov structure, the innovations a_t can be recovered from a properly specified regression using these observations. Assume that the Kalman filter has asymptotically converged, $K_t \rightarrow K$ and $\Sigma_t \rightarrow \Sigma$.

Now we have

$$\begin{aligned}\hat{x}_t &= A_o \hat{x}_{t-1} + K a_{t-1} = A_o \hat{x}_{t-1} + K (y_{t-1} - G \hat{x}_{t-1}) \\ &= (A_o - KG) \hat{x}_{t-1} + K y_{t-1} \\ &= \sum_{j=0}^{\infty} (A_o - KG)^j K y_{t-j-1}\end{aligned}$$

Then

$$y_t = G \hat{x}_t + a_t = G \sum_{j=0}^{\infty} (A_o - KG)^j K y_{t-j-1} + a_t \quad (4.11)$$

Hence, in order to recover the innovations a_t , we need to form an infinite-order autoregression of y_t . The Kalman filter provides an interpretation of the coefficient matrices $(A_o - KG)^j$ estimated in the regression, and the residuals a_t . In practice, the infinite-order VAR is replaced by a VAR with a ‘sufficiently’ large number of lags.

With more structure imposed on the problem, we could be able to identify elements of the innovations vector a_t or their linear combinations with specific structural shocks. Then we can compare impulse responses to that particular innovation generated by the VAR (4.11) with impulse responses to the structural shock in w_t generated by the structural model. [Christiano et al. \(2005\)](#) use this **impulse response function matching** method to find the parameters of the structural model that provide the closest match to a VAR impulse response for a shock to interest rates induced by a monetary policy surprise.

4.2.4 Relationship to linear-quadratic dynamic programming

In Section 3.1.3, we derived the solution to the optimal control problem in a linear-quadratic environment

$$\begin{aligned}v(x_0) &= \max_{\{u_t\}_{t=0}^{\infty}} -E_0 \left[\sum_{t=0}^{\infty} \beta^t \left(x_t' \tilde{R} x_t + u_t' \tilde{Q} u_t \right) \right] \\ x_{t+1} &= A x_t + B u_t + C w_{t+1}\end{aligned}$$

We have determined that the value function is quadratic in the state x_t and optimal control u_t is linear in the state

$$v(x) = -x' P x - \sigma \quad u = -F x \quad (4.12)$$

where the matrices P and F are the fixed points of the value function backward iteration scheme

$$P_{j+1} = \beta (A - B F_j)' P_j (A - B F_j) + \tilde{R} + F_j' Q F_j \quad (4.13)$$

$$F_{j+1} = \beta \left(\tilde{Q} + \beta B' P_j B \right)^{-1} B' P_j A, \quad (4.14)$$

and the evolution of the state under optimal control obtained by solving for the fixed point is given by

$$x_{t+1} = (A - BF)x_t + Cw_{t+1} \quad (4.15)$$

In contrast, the evolution equations for the filtering problem are given by

$$\Sigma_{t+1} = (A_o - K_t G) \Sigma_t (A_o - K_t G)' + CC' + K_t R K_t' \quad (4.16)$$

$$K_t = A_o \Sigma_t G' (G \Sigma_t G' + R)^{-1} \quad (4.17)$$

with evolution of the mean estimate of the state given by

$$\hat{x}_{t+1} = (A_o - K_t G) \hat{x}_t + K_t y_t. \quad (4.18)$$

Asymptotically, matrices Σ_t and K_t will converge to their limiting values Σ and K , respectively, under stationarity assumption imposed on the problem.

Both problems have the same structure, with a reversed time axis. In the linear-quadratic optimization problem, we are exercising optimal control, which leads to the controlled law of motion (4.15), combining the impact of the current state, current control, and the next-period random innovation w_{t+1} . The extent of optimal control is given by the matrix F given in (4.14), which trades off current period cost, summarized by \tilde{Q} , with cost imposed via the impact on the next-period value function, captured by the matrix P_j . Equation (4.13) then summarizes how the value function is iterated backward under the optimally chosen policy to obtain the current iteration P_{j+1} from the previous iteration P_j . The matrix P_{j+1} combines the contribution of P_j , anticipating the effect of the shift in the state under the optimal control, and the current utility flow summarized in $\tilde{R} + F_j' Q F_j$.

On the other hand, the filtering problem yields the evolution of the best forecast of the state (4.18), which combines the impact of the current state, current update via $K_t G \hat{x}_t$, and the new 'shock' y_t . The extent to which the new observation is used in updating the state is controlled by the Kalman gain (4.16), which trades off the impact of the current period noise R and the uncertainty Σ_t inherited from the previous forecast. The evolution equation for Σ_t in (4.16) then summarizes how the covariance of the posterior, which is a quadratic form, evolves over time, combining the uncertainty inherited from the previous period, and the uncertainty embedded in the current period, $CC' + K_t R K_t'$, which combines uncertainty in the evolution of the state, and the uncertainty included in the forecast by embedding the new signal y_t .

The forecasting problem has the same structure as the linear-quadratic dynamic optimization problem because it is a problem that designs the optimal filter under a quadratic objective function. In the linear-Gaussian environment, linear projection that we used minimizes the mean square error of the forecast, which is what is being encoded in the covariance matrix Σ_t . The small difference is that the linear-quadratic decision problem is discounted, which is reflected in the appearance of the time preference parameter β in the equations.

4.3 Estimation

4.3.1 Estimation of the hidden state space model

Since the innovations representation is written completely in terms of a Markov representation of observable variables, we can conveniently write the likelihood function. Recall that we have

$$\begin{aligned}\hat{x}_{t+1} &= A_o \hat{x}_t + K_t a_t \\ y_t &= G \hat{x}_t + a_t\end{aligned}$$

and that $y_t | y^{t-1} \sim N(G\hat{x}_t, \Omega_t)$ where $\Omega_t = \text{Var}(a_t | y^{t-1})$. To see this, notice that $\hat{x}_t = E[x_t | y^{t-1}]$, i.e., \hat{x}_t summarizes information about the mean of x_t embedded in y^{t-1} . Similarly, since $\Sigma_t = \text{Var}(x_t | y^{t-1})$, then $\Omega_t = G\Sigma_t G' + R$. Finally $y_0 \sim N(G\hat{x}_0, \Sigma_0)$.

In this way, the likelihood of observations $\{y_t\}_{t=0}^T$ can be factorized as

$$f(y_0, \dots, y_T) = f(y_T | y^{T-1}) \dots f(y_1 | y^0) f(y_0).$$

Observe the crucial difference to Section 3.1.4—the observed data do not represent a Markov system, hence the conditioning on y^{t-1} . The conditionally Gaussian factors are given by

$$\log f(y_t | y^{t-1}) = -\frac{m}{2} \log(2\pi) - \frac{1}{2} \det(\Omega_t) - \frac{1}{2} a_t' \Omega_t^{-1} a_t.$$

Recall that the matrix Ω_t is constructed from matrices A_o, C, G, R . Now again, we can imagine that these matrices are functions of a parameter vector θ which we need to estimate. Maximizing the likelihood

$$f(y^T | \theta)$$

with respect to θ is the efficient way to proceed.

Observe how we achieved the factorization. We observe the data histories y^{t-1} , from which we construct the best predictor of x_t , summarized by the normal distribution $N(\hat{x}_t, \Sigma_t)$. Then the forecast of y_t conditional on the predictor is the same as the forecast conditional on the whole history y^{t-1} .

One aspect that we have not discussed is the way how the filtered path $\{\hat{x}_t\}_{t=0}^T$ was constructed. Recall that we used a method that used linear regressions to project the evolution of the hidden state on innovations in new data. However, in the linear Gaussian model, the least-squares estimator corresponds to the maximum likelihood estimator.

Bayesian estimation

Often we may want to incorporate some prior knowledge about the parameter θ into the estimation. This can be done by combining the likelihood of the data with the prior using the Bayes theorem.

Specifically, assume that prior to seeing the data, the econometrician has a belief about what θ may be, summarized using the **prior distribution** $\tilde{p}(\theta)$. We are interested in knowing how her belief changes after seeing the data, which serve as evidence about which θ are more probable. By Bayes theorem, the **posterior distribution** for θ then is

$$\tilde{p}(\theta | y^T) = \frac{f(y^T, \theta)}{f(y^T)} = \frac{f(y^T | \theta) \tilde{p}(\theta)}{\int f(y^T | \theta) \tilde{p}(\theta) d\theta}.$$

Observe that now, we are not interested in a single value of θ (as in the case of the maximum likelihood estimator) but in the whole posterior distribution $\tilde{p}(\theta | y^T)$. Statistical theory developed powerful simulation algorithms based on **Monte Carlo Markov chain** (MCMC) methods to construct $\tilde{p}(\theta | y^T)$. With an estimate of the posterior distribution at hand, we can construct posterior statistics of θ that are of our interest, like the posterior mean of θ . The choice of these statistics can be motivated by decision-theoretical considerations that weigh losses associated with incorrect estimates of the parameter vector.

The Bayesian approach can often be computationally more tractable than the maximum likelihood estimation, especially when the likelihood has a complicated structure with many local maxima or almost flat regions, because the posterior distribution can be constructed without resorting to any maximization. Appendix B of Chapter 2 in [Ljungqvist and Sargent \(2012\)](#) describes the construction of the **Bayesian posterior distribution** for θ using Markov chain Monte Carlo methods.

Finally, notice that the Kalman filter has a Bayesian interpretation, with $x_0 \sim N(\hat{x}_0, \Sigma_0)$ representing the prior information about the initial state that can be distinct from the true initial distribution.

Other estimation approaches

There is an important caveat to maximum likelihood being the efficient approach to estimation. The efficiency result is conditional on the assumption that the model specification, i.e., the model structure, distributional assumption imposed on exogenous shocks, etc., is correct. However, the econometrician may doubt the specification, and in the presence of model misspecification, other approaches to estimating underlying parameters may be preferable.

In Section 1.1.6, we introduced the idea of method of moments estimation. [Hansen \(1982\)](#) formalized and generalized this idea in the form of the **generalized method of moments** (GMM) estimator. For example, [Hansen and Singleton \(1983\)](#) use GMM to study implications of a large class of asset pricing models based on moment conditions derived from investors' Euler equations. The GMM estimator only relies on partial identification, i.e., its strength lies in the fact that we do not need to specify a full-fledged equilibrium model. Asset pricing is a particular fruitful area of applying GMM, since finding a well-specified model in asset pricing is particularly hard.

Another stream of estimation techniques, applied by [Christiano et al. \(2005\)](#) and many others, relies on **matching impulse response functions** derived from a structural theoret-

ical model (4.10) and from an empirical VAR model (4.11). The complication arises from the fact that the theoretical model (4.10) can be used to construct impulse responses to structural shocks w_t , while the empirical model involves disturbances in the form of non-structural innovations a_t . The identification problem then consists of mapping the non-structural shocks a_t to their structural counterparts w_t through a suitable rotation of the innovations, so that the impulse responses are comparable.

4.3.2 Frequentist and Bayesian inference

There is also a philosophical divide between the **frequentist view** of the world, captured in the maximum likelihood or moment estimation methods, and the **Bayesian view**, associated with the approach above. The distinction may appear negligible in applications but it is worth taking a note of it. It is worth stressing that regardless the conceptual differences, both approaches have a common goal, to learn about the unknown parameter from observed data.

The **frequentist view**, pioneered by Jerzy Neyman, Egon Pearson, or Abraham Wald, takes the perspective that the unknown parameter takes a particular fixed ‘true’ value θ_0 unknown to the econometrician, and the uncertainty in the world comes in the form of random samples of data \hat{x}^T generated from $P(\theta_\theta)$. The goal is to deal with the randomness in the finite samples of data, perhaps by extending the sample length (up to the whole infinite population). The approach relies on statistical testing (hypothesis testing), in the form of asking whether the data could have been generated from a model with a given $\theta = \theta_0$, as opposed to another $\theta = \theta_1$. *The interest lies in the distribution of data given a value of the parameter.*

The **Bayesian view** follows the tradition of Thomas Bayes, Pierre-Simon Laplace, or Bruno de Finetti. This view takes the data \hat{x}^T as fixed pieces of evidence, while the parameter θ is the source of uncertainty. The decision-maker is endowed with a prior belief about the distribution of the parameter, and the pieces of evidence then sharpen (update) the belief about θ . The outcome is a description of the posterior probability distribution of θ , which tells us which values of θ are more likely than others. *The interest lies in the distribution of the parameter given observed data.*

Chapter 5

Finite difference methods in derivative pricing

Textbook: *Brownian motion and Itô calculus:* Duffie (2001), Chapters 5.A–5.D. Øksendal (2007), Chapters 1–6. *Black–Scholes model:* Duffie (2001), Chapters 5.E–5.H, 6.G–6.I. Øksendal (2007), Chapter 12.3. *Numerical methods:* Judd (1998), Chapter 10, Holmes (2007), Thomas (1995), Candler (2001).

Applications: Merton (1973), Black and Scholes (1973), Cox et al. (1979).

In this chapter, we study the problem of pricing of derivative securities in a continuous-time environment. The problem is as follows. Imagine we have a stock with price Q_t that follows a given process. We are now interested in pricing another financial security whose price is derived from the price of the stock. This derivative security generates a one-time cash flow at time T in the amount $G(Q_T)$, and we are interested in the price of this derivative security at time $t \leq T$.

Black and Scholes (1973) and Merton (1973) provided a path-breaking solution to this problem. The solution is based on an application of the arbitrage pricing theory (APT) of Ross (1976), which states that two assets or portfolios that provide identical payoffs also must have the same price, otherwise investors could create arbitrages by buying the cheaper asset and selling the more expensive asset. We described the concept of an arbitrage in Section 1.6.

The derivative pricing result was formulated in a continuous-time model where uncertainty is driven by a special process called the Brownian motion, and the characterization of the solution takes the form of a partial differential equation (PDE). When this PDE does not have an analytical solution, we must solve it numerically. We focus on a versatile method called the finite difference method, which translate the continuous-time continuous-state space problem into a discrete-time discrete-state space problem on a grid similar to that we used in Section 2.4.2.

In Section 5.1, we introduce the Brownian motion, followed by a concise discussion a class of processes derived from the Brownian motion in Section 5.2. Section 5.3 contains

the formulation of the Black–Scholes model for pricing derivative securities. Section 5.4 presents the finite difference method for solving the associated partial differential equation.

5.1 Brownian motion

The Brownian motion is an important stochastic process that constitutes the foundations of continuous-time stochastic calculus. The term goes back to the Scottish botanist Robert Brown who observed in 1828 the irregular movements of pollen stuck in liquid. Brownian motions are the fundamental building block of stochastic integrals, which are a generalization of Riemann–Stieltjes integrals to a stochastic environment, and lead to a broad class of stochastic processes called Itô processes.

5.1.1 Motivating examples

We can view a Brownian motion as a continuous-time version of the accumulation of a sequence of normally distributed shocks. We first provide a sequence of examples that motivate the intuition underlying the Brownian motion using continuous-time limits of discrete-time models.

Predictable dynamics

A discrete-time deterministic model of capital accumulation can be described by the law of motion for capital

$$k_{t+1} = (1 - \delta_t)k_t + i_t, \quad (5.1)$$

where δ_t is the depreciation rate and i_t is the investment rate. To derive the continuous-time version of the equation, assume a time period of length Δt . Then

$$k_{t+\Delta t} - k_t = i_t \Delta t - \delta_t k_t \Delta t$$

where the terms involving Δt represent investment and depreciation flows. Dividing by Δt and taking the limit as $\Delta t \rightarrow 0$ yields

$$\frac{dk_t}{dt} = i_t - \delta_t k_t.$$

Defining the investment rate $\iota_t = i_t/k_t$, we can rewrite the equation as

$$\frac{dk_t}{dt} \frac{1}{k_t} = \frac{d \log k_t}{dt} = \iota_t - \delta_t,$$

which can be represented as

$$k_t = k_0 \exp \left(\int_0^t (\iota_s - \delta_s) ds \right).$$

We have therefore solved for the stock of capital k_t by integrating up net investment $\iota_s - \delta_s$ along the trajectory of the economy over time on $s \in [0, t]$.

This continuous-time limit expressed in the form of an integral is valid even in situations when functions ι and δ are stochastic, as long as the integral above is valid. In this stochastic case, the dynamics of ι and δ generate a filtration $\{\mathcal{F}_t\}$, $t \in \mathcal{T} = \{0, 1, \dots, T\}$, see Section A.2. For a given path, the integral is a standard Riemann–Stieltjes integral, which follows from the fact that the law of motion (5.1) implies that k_{t+1} is so-called ‘predictable’, i.e., k_{t+1} is \mathcal{F}_t measurable.

Portfolio choice

The predictability assumption used in the preceding example is rather restrictive, and we are interested in constructing continuous-time counterparts of stochastic dynamics where it does not hold. Consider the joint evolution of two investment instruments:

$$\begin{aligned} Q_{t+1} &= Q_t + \mu_t + \sigma_t (W_{t+1} - W_t) \\ B_{t+1} &= B_t + r_t B_t. \end{aligned} \quad (5.2)$$

for $t = 0, 1, \dots, T - 1$. Here, Q_t can be interpreted as the stock price of a non-dividend paying stock, B_t as the cumulative value of investment into a sequence of one-period risk-free bond contracts with one-period interest rate r_t , and $W_{t+1} - W_t \sim N(0, I)$ is a normally distributed shock. The joint dynamics of the two processes generate a filtration $\{\mathcal{F}_t\}$, $t \in \mathcal{T} = \{0, 1, \dots, T\}$. The expected return on the stock is

$$E \left[\frac{Q_{t+1} - Q_t}{Q_t} \mid \mathcal{F}_t \right] = \frac{\mu_t}{Q_t}$$

and σ_t is the one-period volatility of the stock return. At any date t , the investor chooses to invest the current wealth J_t by purchasing θ_t^f units of the risk-free asset at price B_t , and θ_t^r units of the risky asset at price Q_t . The budget constraint is

$$J_t = \theta_t^f B_t + \theta_t^r Q_t.$$

The value of this portfolio at time $t + 1$ is

$$J_{t+1} = \theta_t^f B_{t+1} + \theta_t^r Q_{t+1},$$

which can be subsequently reinvested again. Manipulating this expression yields

$$J_{t+1} = \theta_t^f (B_{t+1} - B_t) + \theta_t^r (Q_{t+1} - Q_t) + \underbrace{\theta_t^f B_t + \theta_t^r Q_t}_{J_t}.$$

Hence we have an expression for the increment in wealth $J_{t+1} - J_t$, which corresponds to the net return on the portfolio. Summing up these increments over time, we have

$$\sum_{t=0}^{T-1} (J_{t+1} - J_t) = J_T - J_0 = \sum_{t=0}^{T-1} \left[\theta_t^f (B_{t+1} - B_t) + \theta_t^r (Q_{t+1} - Q_t) \right].$$

The intertemporal portfolio choice is determined as a solution to the problem of maximizing expected utility from time- T wealth J_T ,

$$E [u (J_T)]$$

subject to the intertemporal budget constraint and initial condition J_0 , with

$$\begin{aligned} J_T &= J_0 + \sum_{t=0}^{T-1} \left[\theta_t^f (B_{t+1} - B_t) + \theta_t^r (Q_{t+1} - Q_t) \right] \\ &= J_0 + \sum_{t=0}^{T-1} \left[\theta_t^f r_t B_t + \theta_t^r \mu_t + \theta_t^r \sigma_t (W_{t+1} - W_t) \right], \end{aligned}$$

where the second line follows from the definition of the returns in (5.2). Repeating the continuous-time approximation above, we now have the dynamics on periods with interval Δt

$$\begin{aligned} Q_{t+\Delta t} - Q_t &= \mu_t \Delta t + \sigma_t (W_{t+\Delta t} - W_t) \\ B_{t+\Delta t} - B_t &= r_t B_t \Delta t, \end{aligned}$$

with $W_{t+\Delta t} - W_t \sim N(0, \Delta t)$. The wealth accumulation process is given by

$$\begin{aligned} J_T &= J_0 + \sum_{i=0}^{I-1} \left[\theta_{i\Delta t}^f (B_{(i+1)\Delta t} - B_{i\Delta t}) + \theta_{i\Delta t}^r (Q_{(i+1)\Delta t} - Q_{i\Delta t}) \right] \\ &= J_0 + \sum_{i=0}^{I-1} \left[\left(\theta_{i\Delta t}^f r_{i\Delta t} B_{i\Delta t} + \theta_{i\Delta t}^r \mu_{i\Delta t} \right) \Delta t + \theta_{i\Delta t}^r \sigma_{i\Delta t} (W_{(i+1)\Delta t} - W_{i\Delta t}) \right] \end{aligned}$$

with $I = T/\Delta t$. We would like to take the continuous-time limit that should lead to

$$\begin{aligned} dQ_t &\approx \mu_t + \sigma_t "dW_t" \\ dB_t &= r_t B_t dt. \end{aligned} \tag{5.3}$$

The question is how to construct the limiting approximation of the stochastic component " dW_t " on the first line rigorously. The limit will lead to a so-called stochastic differential equation, which cannot be characterized by a Lebesgue integral. Uncertainty in Q_t will be driven by innovations to a Brownian motion that could be interpreted as a limiting sequence of normally distributed increments.

An associated problem is the continuous-time limit of the portfolio strategy $\{\theta_t^f, \theta_t^r\}$

that leads to the process for wealth accumulation

$$J_T = J_0 + \int_0^T \left[\left(\theta_t^f r_t B_t + \theta_t^r \mu_t \right) dt + \theta_t^r \sigma_t "dW_t" \right]. \quad (5.4)$$

In the discrete-time model, the investor chooses the portfolio shares θ_t^f, θ_t^r at discrete times $t = 0, 1, \dots, T - 1$, where each pair θ_t^f, θ_t^r is \mathcal{F}_t -measurable. In the continuous-time limit, the investor will adjust the portfolio continuously in a sense that needs to be made precise, and this strategy will be represented by a pair of stochastic processes θ_t^f, θ_t^r , $t \in [0, T]$ that will depend on the observed histories of the shocks, and satisfy certain measurability restrictions.

Filtering

Another example involves linear filtering using the [Kalman \(1960\)](#) filter

$$\begin{aligned} x_{t+1} - x_t &= Ax_t + B(W_{t+1} - W_t) \\ y_{t+1} - y_t &= Dx_t + G(W_{t+1} - W_t) \end{aligned}$$

where y_t represents an m -dimensional vector of observable data, x_t is an n -dimensional vector of an unobservable stochastic process which the decision-maker predicts using observations of y_t , and $W_{t+1} - W_t \sim N(0, I_p)$ is a shock with a p -dimensional multivariate normal distribution. This formulation is only a minor change of notation relative to the version of the filter that we solved in [Section 4.1](#).

We can again rely on the same intuition with discretization

$$\begin{aligned} x_{t+\Delta} - x_t &= Ax_t \Delta t + B(W_{t+\Delta} - W_t) \\ y_{t+\Delta} - y_t &= Dx_t \Delta t + G(W_{t+\Delta} - W_t) \end{aligned}$$

and the limit should again lead to

$$\begin{aligned} dx_t &\approx Ax_t + B "dW_t" \\ dy_t &\approx Dx_t + G "dW_t". \end{aligned}$$

The continuous-time version of the filter was solved by [Kalman and Bucy \(1961\)](#).

The formulation of the portfolio choice and filtering examples seems to be restrictive in the sense that the dynamics of relevant variables, for example the stock return in [\(5.3\)](#) are separated into a predictable component (μ_t) and a component representing the impact of the infinitesimal normally distributed shock (scaled by σ_t). While this induced linearity seems restrictive, and it certainly is restrictive in discrete-time models, this specification provides much more generality in the continuous-time limit.




Figure 5.1: Sample paths of a Brownian motion.

5.1.2 Definition and properties

The definition of the Brownian motion is straightforward.

Definition 5.1. A k -dimensional Brownian motion is a stochastic process W on \mathbb{R}^k such that

1. $W_0 = 0$,
2. $\forall s, t \in \mathcal{T}$ for which $s \leq t$, the difference $W_t - W_s \sim N(0, (t - s) I_k)$,
3. for all $t_0 < t_1 < t_2 < \dots < t_n \in \mathcal{T}$, the random variables $W_{t_j} - W_{t_{j-1}}$, $j \in \{1, \dots, n\}$ are independent.

Said simply, the Brownian motion is a process with independent, normally distributed increments. It turns out that this definition characterizes a unique process, as long as we restrict our attention to processes with continuous sample paths.

Following the general construction of a stochastic process from Section A, we denote Ω the sample space, or the set of all paths of the Brownian motion, with elements $\omega \in \Omega$. The expression $W(\omega)$ represents one particular path of the Brownian motion, and $W_t(\omega)$ the associated value along that path at time t . Figure 5.1 plots simulated sample paths of a Brownian motion.

As any stochastic process, the Brownian motion generates a filtration $\{\mathcal{F}_t\}$ where, somewhat informally, \mathcal{F}_t is the information set that contains all information about the realized path of the Brownian motion up to time t .

The Brownian motion satisfies the **Markov property**: $\forall t, s \geq 0$ and for every Borel set $H \in \mathcal{B}$ on \mathbb{R}^k

$$P(W_{t+s} \in H \mid \mathcal{F}_t) = P(W_{t+s} \in H \mid W_t).$$

See Section A.3 for the definition of Borel sets. In words, the distribution of W_{t+s} conditional on time- t information set is the same as the distribution conditional only on the value W_t . This follows directly from the independence of the increments of the Brownian motion in Definition 5.1.

The Brownian motion is also a **martingale** with respect to filtration $\{\mathcal{F}_t\}$. For $s < t$,

$$E[W_t \mid \mathcal{F}_s] = E[W_t - W_s \mid \mathcal{F}_s] + W_s = W_s$$

and, at the same time,

$$(E[|W_t|])^2 \leq E[|W_t|^2] = nt < +\infty,$$

so that the properties of Definition A.10 are satisfied.

5.1.3 Variation of a stochastic process

We want to establish a formal definition of how ‘variable’ the paths of a stochastic process are. These results will show that the paths of the Brownian motion are qualitatively distinct from the usual smooth functions of time. The idea of the definition is to partition a particular time interval \mathcal{T} , define a discrete-time version of variability of the paths by measuring changes in the value of the stochastic process along the path between the nodes of the partition, and then take a continuous-time limit as the partition is refined and the distance between the nodes of the partition vanishes to zero.

Definition 5.2. *The set of points $\mathcal{P} = \{t_0, \dots, t_n\}$ with $0 = t_0 < t_1 < \dots < t_n = t$ is a **partition of the interval** $[0, t]$. Define*

$$l(\mathcal{P}) = \max |t_j - t_{j-1}|.$$

*to be the **norm of the partition**.*

We denote $l(\mathcal{P}) \rightarrow 0$ to be the limit of an arbitrary sequence of partitions \mathcal{P} such that the norm of the partitions in the sequence converges to zero. Then we can define the following:

Definition 5.3. *Let $X : \Omega \times \mathcal{T} \rightarrow \mathbb{R}$ be a continuous stochastic process. Then for $p > 0$ define the **p -th variation process** of X_t as*

$$\langle X, X \rangle_t^p(\omega) = \lim_{l(\mathcal{P}) \rightarrow 0} \sum_{j=0}^{n-1} |X_{t_{j+1}}(\omega) - X_{t_j}(\omega)|^p$$

where the limit is in probability.

*For $p = 1$, the variation process is called the **total variation process**, and for $p = 2$, it is*

called the **quadratic variation process**. For the second moment, we can also define the **cross-variation process** of two processes X and Y as

$$\begin{aligned}\langle X, Y \rangle_t^2(\omega) &= \lim_{l(\mathcal{P}) \rightarrow 0} \sum_{j=0}^{n-1} \left(X_{t_{j+1}}(\omega) - X_{t_j}(\omega) \right) \left(Y_{t_{j+1}}(\omega) - Y_{t_j}(\omega) \right) \\ &= \frac{1}{4} \left[\langle X + Y, X + Y \rangle_t^2(\omega) - \langle X - Y, X - Y \rangle_t^2(\omega) \right].\end{aligned}\quad (5.5)$$

Example 5.1. (Exercise 2.17 in Øksendal (2007)) For the univariate Brownian motion W ,

$$\langle W, W \rangle_t^2(\omega) = t \quad \text{a.s..}$$

To show this, start with a partition \mathcal{P} of the time interval $[0, t]$, and consider

$$\begin{aligned}E \left[\left(\sum_{t_j \leq t} (W_{t_{j+1}} - W_{t_j})^2 - t \right)^2 \right] &= E \left[\left(\sum_{t_j \leq t} (W_{t_{j+1}} - W_{t_j})^2 \right)^2 \right] - \\ &\quad - 2t \sum_{t_j \leq t} E \left[(W_{t_{j+1}} - W_{t_j})^2 \right] + t^2 \\ &= \sum_{t_j \leq t} 3(t_{j+1} - t_j)^2 + \sum_{\substack{t_j, t_k \leq t \\ j \neq k}} (t_{j+1} - t_j)(t_{k+1} - t_k) - 2t^2 + t^2 = \\ &= 2 \sum_{t_j \leq t} (t_{j+1} - t_j)^2 + t^2 - 2t^2 + t^2 \rightarrow 0\end{aligned}$$

as $l(\mathcal{P}) \rightarrow 0$. Therefore $\sum_{t_j \leq t} (W_{t_{j+1}} - W_{t_j})^2 \rightarrow t$ in $L^2(P)$ sense.

This is a strong result because it shows that every individual path of the Brownian motion on $[0, t]$ has the same ‘length’ t when measured using the quadratic variation. Since $\langle W, W \rangle_{\Delta t}^2(\omega) = \Delta t$ for an arbitrarily short interval Δt , the result also provides heuristic intuition why we can write $(dW_t)^2 = dt$, which is a central insight of Itô calculus, as manifested in Itô’s lemma, Theorem 5.1.

Since the quadratic variation is finite, it can be shown that the **total variation** of a Brownian motion must be infinite,

$$\forall t > 0 : \langle W, W \rangle_t^1(\omega) = +\infty.$$

This conclusion also implies that the paths of a Brownian motion are **nowhere differentiable**.

The Brownian motion is an incredibly versatile process. For example, it turns out that continuous-time martingales with continuous paths and finite quadratic variation can be represented as transformations of a Brownian motion. Informally, the result states that when specifying martingales in a stochastic environment generated by the Brownian motion, such transformed Brownian motions cover the class of martingales without loss of

generality.

5.2 Itô calculus

The Brownian motion is an essential component in the construction of the stochastic integral. The geometric argument underlying the construction is conceptually similar to that of the Riemann–Stieltjes integral but technical complications associated with the irregularity of paths of the Brownian motion are substantial and require a careful treatment.

We start with the review of the construction of the Riemann–Stieltjes integral. In order to construct the Riemann integral of a piecewise continuous function $f(t)$ on $\mathcal{T} = [0, T]$ we choose a partition \mathcal{P} of \mathcal{T} , and then define the integral through the limit

$$\int_0^T f(t) dt \doteq \lim_{l(\mathcal{P}) \rightarrow 0} \sum_{j=0}^{n-1} f(\tau_j) (t_{j+1} - t_j), \quad (5.6)$$

where the values of τ_j can be chosen to be arbitrary elements of the intervals of the partition, $\tau_j \in [t_j, t_{j+1}]$. The integral is well defined only if the limit does not depend on a particular choice of the sequence of partitions, nor on the choices of the points $\tau_j \in [t_j, t_{j+1}]$. Geometrically, the construction approximates the area under the curve $f(t)$ using the sum of rectangular areas $f(\tau_j) (t_{j+1} - t_j)$.

The Riemann–Stieltjes integral generalizes the Riemann integral to integrate along the path of a sufficiently smooth function $g(t)$:

$$\int_0^T f(t) dg(t) \doteq \lim_{l(\mathcal{P}) \rightarrow 0} \sum_{j=0}^{n-1} f(\tau_j) (g(t_{j+1}) - g(t_j)). \quad (5.7)$$

The idea underlying the construction of the stochastic integral is similar, with integration along the path of a smooth function g to be replaced with integration along the path of the Brownian motion $W(\omega)$. In particular, we desire to form the discrete-time approximation using a partition \mathcal{P} ,

$$\sum_{j=0}^{n-1} f(\tau_j) (W_{t_{j+1}}(\omega) - W_{t_j}(\omega)), \quad (5.8)$$

and ask how to construct a well-defined limit as $l(\mathcal{P}) \rightarrow 0$, in the same way the Riemann–Stieltjes integral is formulated in (5.7).

The sum in (5.8) depends on the particular path ω of the Brownian motion, and so it is itself random. Moreover, the integrand $f(\tau_j)$ can also be a stochastic process itself, and so we can write instead $f_{\tau_j}(\omega)$, where f_{τ_j} is a random variable measurable with respect to \mathcal{F}_{τ_j} . The stochastic integral that is the desired outcome of this construction will therefore also be a stochastic process.

It turns out that due to the infinite total variation of W , we will need to choose the points τ_j in a specific way to make the limit well defined. Our particular interest in financial applications will lead us to choose τ_j to be the initial point of the interval, $\tau_j = t_j$, which

yields the so-called **Itô integral**. We proceed in several steps, first providing the definition of the integral for a class of so-called elementary processes, and then extend this definition to larger classes of processes through limits.

Example 5.2. *As a motivation, consider an economically interesting example that builds on the problem introduced in Section 5.1.1. Let the share price evolve as a Brownian motion W , and consider an investor who can trade shares only at a finite number of dates $t_j \in [0, T]$ which define a partition \mathcal{P} .*

Denote $\theta_{t_j}(\omega)$ the number of shares bought at time t_j . We assume that the choice $\theta_{t_j}(\omega)$ can depend on information available up to time t_j , and therefore θ_{t_j} can be viewed as a random variable measurable with respect to \mathcal{F}_{t_j} . With this construction, the evolution of wealth J_t is given by

$$J_t(\omega) = J_0 + \sum_{j=0}^{n(t)-1} \theta_{t_j}(\omega) \left(W_{t_{j+1}}(\omega) - W_{t_j}(\omega) \right) + \theta_{t_{n(t)}}(\omega) \left(W_t(\omega) - W_{t_{n(t)}}(\omega) \right) \quad (5.9)$$

where $n(t)$ is the index of the interval in the partition such that $t \in [t_{n(t)}, t_{n(t)+1})$, and $n(T) = n$. The wealth process represents cumulative gains from investments between individual trading dates, and the last term is the gain until time t since the last preceding trading date $t_{n(t)}$. The key observation is the timing procedure, which makes the portfolio choice θ_{t_j} for the investment in the time interval $[t_j, t_{j+1})$ measurable with respect to the σ -algebra \mathcal{F}_{t_j} . The process θ_t viewed as a continuous-time process is constant on the intervals $[t_j, t_{j+1})$, and is called a **dynamic strategy**.

5.2.1 Construction of Itô integral

Processes that have piecewise constant trajectories that are allowed to jump only at a finite number of times are called elementary processes.

Definition 5.4. *An elementary (also called simple) process ϕ on $[0, T]$ is a process for which there exists a partition \mathcal{P} of $[0, T]$ such that $\phi_t = \phi_{t_j}$ for $t \in [t_j, t_{j+1})$.*

Stochastic integrals of elementary processes constitute the fundamental building blocks in our analysis, and they are defined via a direct construction. The construction exactly mimics the intuition of the dynamic strategy outlined in (5.9).

Definition 5.5. *For the class of elementary processes ϕ , the Itô integral of ϕ is defined as*

$$\int_0^t \phi_s(\omega) dW_s(\omega) \doteq \sum_{j=0}^{n(t)-1} \phi_{t_j}(\omega) \left(W_{t_{j+1}}(\omega) - W_{t_j}(\omega) \right) + \phi_{t_{n(t)}}(\omega) \left(W_t(\omega) - W_{t_{n(t)}}(\omega) \right), \quad (5.10)$$

where the last term reflects the interrupted last interval of the partition and $n(t)$ is such that $t \in [t_{n(t)}, t_{n(t)+1})$.

Since the elementary process is constant on $[t_{n(t)}, t_{n(t)+1})$, we could have chosen the value ϕ_{τ_j} for any $\tau_j \in [t_j, t_{j+1})$ without any consequences. The integral in (5.10) is itself a

stochastic process, since the integral is defined pathwise, for every $\omega \in \Omega$. The argument ω will mostly be omitted in the expressions in the following text.

Stochastic integrals of more complicated processes will be constructed through limiting arguments, using approximations of integrated processes on the class elementary processes. The following example illustrates the challenges in such an approximation.

Example 5.3. *We are interested in approximating the evolution of a Brownian motion W on the interval $\mathcal{T} = [0, T]$. Consider a partition \mathcal{P} of \mathcal{T} , and define two candidate approximations,*

$$\begin{aligned}\phi_t(\omega) &= \sum_{j=0}^{n-1} W_{t_j}(\omega) \mathbf{1}_{[t_j, t_{j+1})}(t) \\ \psi_t(\omega) &= \sum_{j=0}^{n-1} W_{t_{j+1}}(\omega) \mathbf{1}_{[t_j, t_{j+1})}(t).\end{aligned}$$

The process ϕ_t approximates the path of the Brownian motion $W_t(\omega)$ on the interval $[t_j, t_{j+1})$ with the initial value $W_{t_j}(\omega)$, while ψ_t approximates the path of the Brownian motion $W_t(\omega)$ on the interval $[t_j, t_{j+1})$ with the terminal value $W_{t_{j+1}}(\omega)$. The values at the endpoints of the intervals are inconsequential. As we refine the partition, the approximations therefore approach each other in a loose sense, so it would seem that choosing ϕ_t or ψ_t will lead to the same limit as $l(\mathcal{P}) \rightarrow 0$.

In the case of a Riemann integral, the limits of integrals of the two approximations as we refine the partition indeed coincide, and define the Riemann integral of the Brownian motion:

$$\lim_{l(\mathcal{P}) \rightarrow 0} \int_0^T \phi_t(\omega) dt = \lim_{l(\mathcal{P}) \rightarrow 0} \int_0^T \psi_t(\omega) dt \doteq \int_0^T W_t(\omega) dt.$$

This is not surprising because this construction is in line with the definition of Riemann integral, and since the path of a Brownian motion is continuous, the path is Riemann integrable.

The situation is markedly different in the case of the stochastic integral. Since ϕ and ψ are elementary processes, their Itô integrals on $[0, T]$ are defined, omitting the path arguments ω , as

$$\begin{aligned}\int_0^t \phi_s dW_s &\doteq \sum_{j=0}^{n(t)-1} W_{t_j} (W_{t_{j+1}} - W_{t_j}) + W_{t_{n(t)}} (W_t - W_{t_{n(t)}}) \\ \int_0^t \psi_s dW_s &\doteq \sum_{j=0}^{n(t)-1} W_{t_{j+1}} (W_{t_{j+1}} - W_{t_j}) + W_{t_{n(t)+1}} (W_t - W_{t_{n(t)}}).\end{aligned}$$

To see the distinction between the two constructions, compute the expectations

$$\begin{aligned}E \left[\int_0^T \phi_t dW_t \mid \mathcal{F}_0 \right] &= E \left[\sum_{j=0}^{n-1} W_{t_j} (W_{t_{j+1}} - W_{t_j}) \mid \mathcal{F}_0 \right] = 0, \\ E \left[\int_0^T \psi_t dW_t \mid \mathcal{F}_0 \right] &= E \left[\sum_{j=0}^{n-1} W_{t_{j+1}} (W_{t_{j+1}} - W_{t_j}) \mid \mathcal{F}_0 \right] = T.\end{aligned}\tag{5.11}$$

These expectations will therefore remain distinct even if we take the limit $l(\mathcal{P}) \rightarrow 0$. So despite the fact that both ϕ and ψ seem to be reasonable approximations of W for the purposes of the construction the stochastic integral, they yield different answers.

The intuitive reason is that there is too much variation in W over time—the result is closely related to the fact that W is a process of infinite total variation, while the Stieltjes integral (5.7) integrates against a function g that has finite total variation. Hence approximations in the construction of the stochastic integral using values at the initial or terminal endpoint of the interval do not yield results that are close to each other as we refine the partition.

Example 5.3 shows that the choice of the points at which to approximate the integrated function on the subintervals of the partition is essential. The choice of the approximation of W using ϕ , i.e., at the initial points of the subintervals, yields the **Itô integral**. Other choices are possible as well, for example, choosing the average $\frac{1}{2}(\phi + \psi)$ as the approximation yields the so-called **Stratonovich integral**.

From the perspective of financial applications, it turns out that the Itô integral is the desirable choice. This choice aligns with the wealth accumulation process in Example 5.2, the resulting Itô integral of the portfolio choice process θ yields the appropriate wealth dynamics, with a portfolio choice process θ that corresponds to infinitesimal rebalancing as $l(\mathcal{P}) \rightarrow 0$. Mathematically, the Itô integral is adapted to filtration $\{\mathcal{F}_t\}$ generated by the Brownian motion, meaning that the portfolio choice at the t cannot depend on information at future dates $u > t$.

On the other hand, the integral constructed using ψ is not adapted to $\{\mathcal{F}_t\}$ because the integral up to time $t \in [t_j, t_{j+1})$ uses the value of ψ at $t_{j+1} > t$. While this may seem innocuous in the limit as $l(\mathcal{P}) \rightarrow 0$, it is a key difference. In the wealth accumulation example, this corresponds to choosing the portfolio θ_{t_j} at time t_j according to the realized return $W_{t_{j+1}}$ at the future date t_{j+1} .

The Itô integral is also a martingale with respect to the filtration generated by the Brownian motion, consistently with the expected value in (5.11). For $u < t$,

$$\begin{aligned} E \left[\int_0^t \phi_s dW_s \mid \mathcal{F}_u \right] &= E \left[\int_0^u \phi_s dW_s + \int_u^t \phi_s dW_s \mid \mathcal{F}_u \right] \\ &= \int_0^u \phi_s dW_s + E \left[\int_u^t \phi_s dW_s \mid \mathcal{F}_u \right] = \int_0^u \phi_s dW_s. \end{aligned} \quad (5.12)$$

The first equality follows from the linearity of the integral, the second equality from the fact that $\int_0^u \phi_s dW_s$ is measurable with respect to \mathcal{F}_u , and the third equality utilizes the same calculation as in (5.9).

Having defined the Itô integral on the class of elementary processes in Definition 5.5, the goal now is to extend this definition to a larger class of stochastic processes f , giving a meaning to the expression

$$\int_0^t f_s(\omega) dW_s(\omega).$$

The construction is based on limiting approximations of the stochastic process f using elementary processes. Naturally, for the construction to work, we need to make sure that the

class of stochastic processes of interest can be well approximated by elementary processes. We will not provide the details of this construction, which can be found, for example, in Øksendal (2007), Chapter 3. We will however point out some results.

Let \mathcal{L} be the set of all processes adapted to the filtration generated by the Brownian motion. We define three nested subsets of this set:

$$\mathcal{L}^1 = \left\{ f \in \mathcal{L} : \int_0^T |f_t| dt < \infty \text{ a.s.} \right\} \quad (5.13)$$

$$\mathcal{L}^2 = \left\{ f \in \mathcal{L} : \int_0^T (f_t)^2 dt < \infty \text{ a.s.} \right\} \quad (5.14)$$

$$\mathcal{H}^2 = \left\{ f \in \mathcal{L}^2 : E \left[\int_0^T (f_t)^2 dt \right] < \infty \right\} \quad (5.15)$$

The class \mathcal{L}^1 is the class of all pathwise integrable processes, \mathcal{L}^2 is the class of all pathwise square integrable processes, and \mathcal{H}^2 the class of processes with a finite second moment. It is easy to see that $\mathcal{H}^2 \subset \mathcal{L}^2 \subset \mathcal{L}^1$.

The choice of the classes is motivated by the desire to construct stochastic integrals for a sufficiently large set of functions that will be of our interest. Specifically, for stochastic processes $f \in \mathcal{H}^2$, we have the following definition.

Definition 5.6. Let $f \in \mathcal{H}^2$. Then the *Itô integral* of f is defined by

$$\int_0^T f_t(\omega) dW_t(\omega) = \lim_{k \rightarrow \infty} \int_0^T \phi_t^k(\omega) dW_t(\omega) \quad (5.16)$$

where $\{\phi^k\}$ is a sequence of elementary functions in \mathcal{H}^2 such that

$$E \left[\int_0^T \left(f_t(\omega) - \phi_t^k(\omega) \right)^2 dt \right] \rightarrow 0 \text{ as } k \rightarrow \infty. \quad (5.17)$$

The idea of the definition is to construct a sequence of elementary processes $\{\phi^k\}$ with piecewise constant paths such that ϕ^k approximates the process f successively better as $k \rightarrow \infty$, in the sense of (5.17). The definition is only meaningful if every such sequence $\{\phi^k\}$ that satisfies (5.17) yields the same value of the limit on the right-hand side of (5.16), which is a result that needs to be proven. Then this common value defines the Itô integral of f .

Functions in the class \mathcal{H}^2 are well behaved, and the stochastic integral of a process $f \in \mathcal{H}^2$ has the martingale property, just like in the case of elementary processes (5.12):

$$E \left[\int_0^t f_s(\omega) dW_s(\omega) \mid \mathcal{F}_u \right] = \int_0^u f_s(\omega) dW_s(\omega).$$

Stochastic integrals can also be defined in a similar way for processes f belonging to the larger class \mathcal{L}^2 but the resulting stochastic integrals may not be martingales anymore, and are generally only so-called local martingales.

5.2.2 Itô processes

We now define a class of processes called Itô processes that additively combine Riemann integrals and Itô integrals. As it will turn out, this class is sufficiently general to cover many interesting applications.

Definition 5.7. *An n -dimensional Itô process is a process $X : \Omega \times \mathcal{T} \rightarrow \mathbb{R}^n$ such that*

$$X_t = X_0 + \int_0^t \mu_s ds + \int_0^t \sigma_s dW_s \quad (5.18)$$

where $\mu \in (\mathcal{L}^1)^n$, $\sigma \in (\mathcal{L}^2)^{n \times k}$ and W is a k -dimensional Brownian motion. We assume that μ and σ are \mathcal{F}_t -adapted where $\{\mathcal{F}_t\}$ is a filtration with respect to which W is a martingale.

An **Itô diffusion** is an Itô process for which the coefficients satisfy $\mu_s = \mu(X_s)$ and $\sigma_s = \sigma(X_s)$ for all $s \in \mathcal{T}$.

The process μ is called **drift**, and σ is called the **volatility** of the Itô process. The Itô process X_t defined above is n -dimensional, with uncertainty generated by a k -dimensional Brownian motion, hence μ is an $n \times 1$ -dimensional vector process, and σ is an $n \times k$ dimensional. The restrictions $\mu \in (\mathcal{L}^1)^n$ and $\sigma \in (\mathcal{L}^2)^{n \times k}$ are imposed so that the Riemann integral and the Itô integral in (5.18) are well defined. The definition assures that $\int_0^t \sigma_s dW_s$ is at least a local martingale, and when $\sigma \in (\mathcal{H}^2)^{n \times k}$, then it is also a martingale.

As stated above, an Itô process is the sum of a Riemann integral and an Itô integral. Often, equation (5.18) is written in the 'differential' form

$$dX_t = \mu_t dt + \sigma_t dW_t.$$

When the processes $\mu, \sigma \in \mathcal{H}^2$, then the Itô integral is a martingale and the argument of the Itô integral is square integrable. It then follows that for $t, u \geq 0$,

$$\begin{aligned} E[X_{t+u} | \mathcal{F}_t] &= X_t + \int_t^{t+u} \mu_s ds \\ \text{Var}[X_{t+u} | \mathcal{F}_t] &= E \left[\left(\int_t^{t+u} \sigma_s dW_s \right)^2 \middle| \mathcal{F}_t \right] = E \left[\int_t^{t+u} |\sigma_s|^2 dt \middle| \mathcal{F}_t \right] \end{aligned}$$

where the last equality follows from a result known as Itô isometry. Then we can localize the mean and variance by constructing the infinitesimal expected growth rate and variance of the process:

$$\begin{aligned} \frac{d}{du} E[X_{t+u} | \mathcal{F}_t] \Big|_{u=0} &= \mu_t \quad \text{a.s.} \\ \frac{d}{du} \text{Var}[X_{t+u} | \mathcal{F}_t] \Big|_{u=0} &= |\sigma_t|^2 = \sigma_t \sigma_t' \quad \text{a.s.} \end{aligned}$$

which provides a justification for calling the two coefficients the drift and volatility of the

Itô process. Informally, we will write these results in the shorthand differential form

$$\begin{aligned} E_t [dX_t] &= \mu_t dt \\ \text{Var}_t [dX_t] &= \sigma_t \sigma_t' dt, \end{aligned}$$

where $E_t [\cdot] = E [\cdot | \mathcal{F}_t]$.

Given an Itô process X , we define two additional classes of processes that will be useful in the study of portfolio problems and other applications.

Definition 5.8. For an n -dimensional Itô process X , we define

$$\mathcal{L}(X) = \left\{ \theta \in (\mathcal{L})^n : \theta' \mu \in \mathcal{L}^1, \theta' \sigma \in (\mathcal{L}^2)^k \right\} \quad (5.19)$$

$$\mathcal{H}^2(X) = \left\{ \theta \in \mathcal{L}(X) : E \left[\left(\int_0^T \theta_t' \mu_t dt \right)^2 \right] < \infty, \theta' \sigma \in (\mathcal{H}^2)^k \right\} \quad (5.20)$$

Definition 5.8 can be interpreted as follows. Let X be the vector process describing the price dynamics of n available assets with drift μ and volatility σ , and θ is an $n \times 1$ dynamic portfolio strategy with each entry representing the number of units of each of these assets held over time. Then the instantaneous gain on this portfolio at time t is $\theta_t' dX_t$, and accumulating these gains over time yields

$$J_T = J_0 + \int_0^T \theta_t' dX_t = J_0 + \int_0^T \theta_t' \mu_t dt + \int_0^T \theta_t' \sigma_t dW_t, \quad (5.21)$$

which is the wealth process that characterizes the cumulative value of the portfolio gains. Classes $\mathcal{L}(X)$ and $\mathcal{H}^2(X)$ are therefore classes of portfolio strategies θ that leads to portfolio value processes with desirable properties. In particular, if $\theta \in \mathcal{H}^2(X)$, then the stochastic integral $\int_0^T \theta_t' dX_t$ has finite variance and the Itô integral component $\int_0^t \theta_s' \sigma_s dW_s$ is a martingale.

The wealth accumulation process (5.21) is a generalization and the continuous-time version of the wealth accumulation process specified in the motivating example in (5.4). Investor's portfolio strategies are adjusted continuously, in response to changes in economic environment. Since the portfolio strategy consists of continuously reinvesting existing wealth, it is called self-financing.

5.2.3 Itô's lemma

The definition of an Itô process X in (5.18) seems to be restrictive, since it involves a linear combination of a Riemann integral and an Itô integral. It would then seem that nonlinear transformations of X would not longer be Itô processes. For example, in the case of the discrete-time linear vector-autoregression (3.1):

$$x_{t+1} = A_0 x_t + C w_{t+1} \quad w_{t+1} \sim N(0, I_p) \text{ iid},$$

a transformation $y_t = f(x_t)$ for some nonlinear function f would no longer yield a linear vector-autoregression for y_t .

Starting from a given Itô process X , we want to characterize its nonlinear transformation $Y_t = f(t, X_t)$ where f is a given, sufficiently differentiable function. It turns out that Y_t is again an Itô process.

The characterization is provided by a key result of stochastic calculus, Itô's lemma, due to Itô (1951). We first provide its scalar version, with only a heuristic proof.

Theorem 5.1 (Itô's lemma). *Let X be a univariate Itô process*

$$dX_t = \mu_t dt + \sigma_t dW_t$$

where W is a univariate Brownian motion. Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ with $f \in C^2(\mathcal{T} \times \mathbb{R})$ (twice continuously differentiable). Then $Y_t = f(t, X_t)$ is an Itô process and

$$dY_t = f_t(t, X_t) dt + f_x(t, X_t) \mu_t dt + \frac{1}{2} f_{xx}(t, X_t) \sigma_t^2 dt + f_x(t, X_t) \sigma_t dW_t.$$

Proof. The heuristic proof goes as follows. First consider a 'second-order' Taylor approximation

$$dY_t = f_t dt + f_x dX_t + \frac{1}{2} f_{tt} (dt)^2 + f_{tx} dt dX_t + \frac{1}{2} f_{xx} (dX_t)^2$$

Now observe

$$\begin{aligned} dt dX_t &= dt (\mu_t dt + \sigma_t dW_t) = \mu_t (dt)^2 + \sigma_t dt dW_t \\ (dX_t)^2 &= \mu_t^2 (dt)^2 + 2\mu_t \sigma_t dt dW_t + \sigma_t^2 (dW_t)^2 \end{aligned}$$

When we computed the quadratic variation of an Itô process in Example 5.1, we already argued that $(dW_t)^2 = dt$, hence $(dW_t)^2$ is a first-order term in dt . However, the remaining terms are higher than first order. Since dW_t can be argued to have mean zero and variance dt , the term $dt dW_t$ will be mean zero and variance $(dt)^2$, which is a higher-order stochastic term than dW_t . Therefore, the only term left in the two expressions above is $\sigma_t^2 (dW_t)^2 = \sigma_t^2 dt$. Combining these results yields the statement of Itô's lemma. ■

A key observation obtained from Itô's lemma is that the process Y_t also follows an Itô diffusion:

$$Y_t = Y_0 + \int_0^t \left[f_t(s, X_s) + f_x(s, X_s) \mu_s + \frac{1}{2} f_{xx}(s, X_s) \sigma_s^2 \right] ds + \int_0^t f_x(s, X_s) \sigma_s dW_s.$$

The linearity of the Itô process and additivity of its two integrals is therefore without loss of generality, and preserved under the nonlinear transformation $Y_t = f(t, X_t)$. The non-linearity is embedded in the transformation of the drift and volatility coefficients of the Itô process.

The formula can be extended to multivariate Brownian motions when we note that for two independent Brownian motions W^j and W^k , we have $(dW_t^j)(dW_t^k) = 0$.

Theorem 5.2 (Multivariate Itô's lemma). *Let W be a k -dimensional Brownian motion, X an n -dimensional Itô process*

$$dX_t = \mu_t dt + \sigma_t dW_t$$

and $f : \mathcal{T} \times \mathbb{R}^n \rightarrow \mathbb{R}^m$ be from C^2 . Then for $Y_t = f(t, X_t)$, we have for the k -th component Y_t^k

$$dY_t^k = \left[f_t^k + f_x^k \mu_t + \frac{1}{2} \text{tr} \left[\sigma_t \sigma_t' f_{xx}^k \right] \right] dt + f_x^k \sigma_t dW_t.$$

5.2.4 Examples

A useful result in integration is the integration by parts formula for the Riemann–Stieltjes integral. The Itô integral also has an integration by parts formula but only for integrated functions that have bounded total variation, in the sense of Definition 5.3.

Theorem 5.3 (Integration by parts). *Suppose the process $f_t(\omega)$ is continuous and of bounded total variation with respect to $t \in [0, T]$ for almost all ω . Then*

$$\int_0^T f_t dW_t = f_T W_T - \int_0^T W_t df_t. \quad (5.22)$$

The assumption of bounded variation of f is crucial. The integration by parts formula will, for instance, not work for $f_t = W_t$, since the Brownian motion has infinite total variation. The following example illustrates this.

Example 5.4. *We want to show that*

$$\int_0^T W_t dW_t = \frac{1}{2} W_T^2 - \frac{1}{2} T.$$

We can obtain this result as an application of Itô's lemma. Let $X_t = W_t$ and $Y_t = f(t, X_t) = \frac{1}{2} W_t^2$. Then an application of Itô's lemma yields

$$dY_t = d\left(\frac{1}{2} W_t^2\right) = W_t dW_t + \frac{1}{2} dt.$$

and integrating up, we obtain:

$$\begin{aligned} \int_0^T dY_t &= Y_T - Y_0 = \int_0^T W_t dW_t + \int_0^T \frac{1}{2} dt \\ \frac{1}{2} W_T^2 &= \int_0^T W_t dW_t + \frac{1}{2} T. \end{aligned}$$

Notice that if we had used the integration by parts formula (5.22), we would have obtained

$$\int_0^T W_t dW_t \stackrel{!}{=} W_T^2 - \int_0^T W_t dW_t,$$

and hence

$$\int_0^T W_t dW_t = \frac{1}{2} W_T^2,$$

which is an incorrect result because the integration by parts formula omitted the second-order term implied by Itô's lemma.

Example 5.5. We want to compute $\int_0^T t dW_t$. Since t has finite total variation, we can apply the integration by parts formula (5.22)

$$\int_0^T t dW_t = TW_T - \int_0^T W_t dt.$$

This can be verified by applying Itô's lemma to the first term on the right-hand side

$$d(tW_t) = W_t dt + t dW_t.$$

Integrating up verifies that the application of the integration by parts formula is correct in this case.

Itô's lemma can also be useful for the derivation of explicit solutions for Itô's processes, as we illustrate with the next couple of examples.

Example 5.6. Let X be an Itô process characterized in differential form by

$$dX_t = \mu dt + \sigma dW_t$$

with a given initial condition X_0 . We can proceed by integrating

$$\begin{aligned} \int_0^t dX_s &= X_t - X_0 \\ &= \int_0^t \mu ds + \int_0^t \sigma dW_s = \mu \int_0^t ds + \sigma \int_0^t dW_s = \mu t + \sigma (W_t - W_0). \end{aligned}$$

Since $W_0 = 0$, we obtain the explicit solution for X_t in the form

$$X_t = X_0 + \mu t + \sigma W_t,$$

which is a process called the arithmetic Brownian motion. In particular, since $W_t \sim N(0, t)$, the distribution of X_t conditional on X_0 is

$$X_t \sim N(X_0 + \mu t, \sigma^2 t).$$

Example 5.7. Let X be an Itô process characterized in differential form by

$$dX_t = \mu X_t dt + \sigma X_t dW_t$$

with a given initial condition X_0 . To find an explicit solution for X_t , we cannot integrate both sides of the above formulas since the right-hand side also depends on X_t . Let us therefore first define

$Y_t = \log X_t$ and apply Itô's lemma

$$\begin{aligned} dY_t &= d \log X_t = \frac{1}{X_t} dX_t - \frac{1}{2} \frac{1}{X_t^2} (dX_t)^2 = \frac{1}{X_t} (\mu X_t dt + \sigma X_t dW_t) - \frac{1}{2} \frac{1}{X_t^2} \sigma^2 X_t^2 dt \\ &= \left(\mu - \frac{1}{2} \sigma^2 \right) dt + \sigma dW_t. \end{aligned}$$

We can now integrate both sides of the equation

$$\int_0^t dY_s = Y_t - Y_0 = \int_0^t \left(\mu - \frac{1}{2} \sigma^2 \right) ds + \int_0^t \sigma dW_s = \left(\mu - \frac{1}{2} \sigma^2 \right) t + \sigma W_t.$$

Hence we obtain

$$Y_t = Y_0 + \left(\mu - \frac{1}{2} \sigma^2 \right) t + \sigma W_t.$$

Exponentiating and noticing that $X_t = \exp(Y_t)$, we have the solution

$$X_t = X_0 \exp \left(\left(\mu - \frac{1}{2} \sigma^2 \right) t + \sigma W_t \right),$$

which is a process called the geometric Brownian motion. X_t is therefore a random variable that is log-normally distributed conditional on X_0 ,

$$X_t \sim N \left(\log X_0 + \left(\mu - \frac{1}{2} \sigma^2 \right) t, \sigma^2 t \right),$$

and

$$E[X_t | X_0] = X_0 \exp(\mu t),$$

which follows from the formula for the mean of a log-normally distributed variable.

5.2.5 Feynman–Kac formula

The Feynman–Kac formula establishes a link between a class of partial differential equations (PDEs) and stochastic processes driven by Brownian motions. While [Kac \(1949\)](#) derived the result in the diffusion setup, Richard Feynman previously characterized analogous results in the context of quantum mechanics in his thesis.

The formula allows to solve PDEs using simulations of stochastic processes, or, vice versa, solve stochastic differential equations by characterizing solutions to PDEs.

For notational simplicity, we outline the results in the case of a scalar state variable x , the conceptually identical multivariate case can be found, for example, in [Øksendal \(2007\)](#), Theorem 8.2.1. Consider the PDE

$$h(x, t) - g(x, t) r(x, t) + g_x(x, t) \mu(x, t) + \frac{1}{2} g_{xx}(x, t) \sigma(x, t)^2 + g_t(x, t) = 0 \quad (5.23)$$

with terminal boundary condition $g(x, T) = G(x)$. The Feynman–Kac formula states that

the solution $g(x, t)$ to the PDE can be written as a conditional expectation

$$g(x, t) = E \left[\int_t^T \phi(t, s) h(X_s, s) ds + \phi(t, T) G(X_T) \mid X_t = x \right] \quad (5.24)$$

where

$$\phi(t, s) = \exp \left(- \int_t^s r(X_\tau, \tau) d\tau \right)$$

for a stochastic process X that satisfies, under the probability measure under which the expectation is taken,

$$dX_t = \mu(X_t, t) dt + \sigma(X_t, t) dW_t$$

and $X_t = x$ is the initial condition. The solution therefore takes the form of a present value equation, with a payoff flow $h(x, s)$ and terminal payoff $G(x)$, discounted by the discount rate $r(x, t)$.

To show this relationship, we can write equation (5.24) in a recursive form as

$$\begin{aligned} g(x, t) &= h(x, t) dt + E \left[\int_{t+dt}^T \phi(t, s) h(X_s, s) ds + \phi(t, T) G(X_T) \mid X_t = x \right] = \\ &= h(x, t) dt + \phi(t, t+dt) E [g(X_{t+dt}, t+dt) \mid X_t = x] \end{aligned}$$

An application of Itô's lemma to the term $\phi(t, dt) g(X_{t+dt}, t+dt)$ yields

$$\begin{aligned} \phi(t, t+dt) g(X_{t+dt}, t+dt) &= g(X_t, t) - r(X_t, t) g(X_t, t) dt \\ &\quad + \left[g_x(X_t, t) \mu(X_t, t) + \frac{1}{2} g_{xx}(X_t, t) \sigma(X_t, t)^2 + g_t(X_t, t) \right] dt \\ &\quad + g_x(X_t, t) \sigma(X_t, t) dW_t \end{aligned}$$

This expansion can be plugged back into the expectation operator. The contribution of the diffusion term on the last row to the expectation is zero. Reorganizing the terms and dividing the whole equation by dt yields PDE (5.23). The terminal condition on the PDE is satisfied.

5.3 The Black–Scholes model

The Black–Scholes model for option pricing has been developed in [Black and Scholes \(1973\)](#), with a central insight based on dynamic hedging provided by Robert Merton. The model was extended to the pricing of more complicated derivative securities in [Merton \(1973\)](#), and to more complex environments in the subsequent literature. While research on the pricing of derivative securities has been active before, the central contribution of [Black and Scholes \(1973\)](#) and [Merton \(1973\)](#) is that they were able to derive the valuation formulas in terms of relatively easy-to-measure parameters.

5.3.1 Market structure

The environment is a security market model with two basic securities. Time is continuous and given by a finite interval $\mathcal{T} = [0, T]$. One security provides a constant risk-free return r over each infinitesimal horizon. We call this security the (infinitesimal) risk-free bond. An initial investment $B_0 = 1$ into this security accumulates over time as

$$dB_t = rB_t dt \quad (5.25)$$

so that the value of such an investment at time t is

$$B_t = \exp\left(\int_0^t r ds\right) = \exp(rt). \quad (5.26)$$

The second security is a risky stock with price Q_t that follows a geometric Brownian motion

$$dQ_t = \mu Q_t dt + \sigma Q_t dW_t \quad (5.27)$$

with constant scalar parameters μ and σ and a given initial price Q_0 . For simplicity, we assume that the stock pays no dividends on $[0, T]$. An explicit solution to this equation is found in the same way as in Example 5.7. Applying Itô's lemma, we have

$$d \log Q_t = \left(\mu - \frac{1}{2}\sigma^2\right) dt + \sigma dW_t$$

and integrating up, we obtain

$$Q_t = Q_0 \exp\left(\left(\mu - \frac{1}{2}\sigma^2\right)t + \sigma W_t\right). \quad (5.28)$$

The expected price conditional on Q_0 then is

$$E[Q_t|Q_0] = Q_0 \exp(\mu t).$$

From these expressions, we can compute the annualized expected infinitesimal returns on the two investments over a short horizon t . For the investment into the risk-free security

$$\lim_{t \rightarrow 0} \frac{1}{t} \frac{E[B_t] - B_0}{B_0} = \lim_{t \rightarrow 0} \frac{1}{t} (\exp(rt) - 1) = r,$$

and for the risky stock

$$\lim_{t \rightarrow 0} \frac{1}{t} \frac{E[Q_t] - Q_0}{Q_0} = \lim_{t \rightarrow 0} \frac{1}{t} (\exp(\mu t) - 1) = \mu.$$

The infinitesimal risk premium on the stock therefore is $\mu - r$.

The security market is hence characterized by three parameters, the risk-free rate r , the expected return on the risky investment μ and the volatility of the risky investment σ .

Uncertainty in this financial market is driven by a scalar Brownian motion. At every time t , an investor can choose to purchase θ_t^f units of the risk-free asset at price B_t and θ_t^r units of the stock at price Q_t . The financial gain over an infinitesimal horizon from this investment is

$$\theta_t^f dB_t + \theta_t^r dQ_t$$

and a given portfolio strategy θ^f, θ^r yields terminal wealth at time T

$$J_T = J_0 + \int_0^T \left[\left(\theta_t^f r B_t + \theta_t^r \mu \right) dt + \theta_t^r \sigma dW_t \right].$$

The value J_T is the terminal payoff from the portfolio strategy.

We hence have a market with uncertainty driven by a univariate Brownian motion and two investable assets, one risk-free and one risky with a nontrivial volatility of the price. Without going into formal proofs, this market is so-called **complete**. Consider an arbitrary time- T payoff G_T that is \mathcal{F}_T -measurable, which means that it is a function of the history of the Brownian motion W , and, for technical reasons, also square integrable. Market completeness means that any such payoff G_T can be obtained as an outcome of a suitably chosen dynamic portfolio strategy θ^f, θ^r , with a particular amount of initial wealth.

5.3.2 Derivative securities

We are interested in pricing a security with terminal payoff at time T equal to $G(Q_T)$. Since the payoff is a function of the underlying stock price, such a security is called **derivative**. Typical examples of derivative securities are options. A call option with strike price K has payoff

$$G(Q_T) = \max(Q_T - K, 0) \equiv (Q_T - K)_+, \quad (5.29)$$

while a put option with strike price K has payoff

$$G(Q_T) = \max(K - Q_T, 0) \equiv (K - Q_T)_+. \quad (5.30)$$

The term option comes from the fact that, for example in the case of a call option, its payoff is equivalent to the right to buy the underlying stock at time T for the price K . The value of that option at maturity is therefore equal to the difference $Q_T - K$ but if this difference is negative, the owner is not mandated to buy the stock at price K that would be higher than the current market price Q_T . Correspondingly, a put option gives the owner the right to sell the stock at time T for the price K . Options with this specification are called European options because they give the right to exercise them on a given date T , while American options give the right to exercise at any time on or before the given date T .

The time- T payoff $G(Q_T)$ of the derivative security only depends on Q_T . Further, the interest rate r is constant, and the distribution of the future stock price conditional on time- t information only depends on Q_t . We can therefore conjecture that the time- t price can be written as $g(Q_t, t)$, where g is a pricing function we need to derive. The explicit dependence of the price on time t is given by the fact that we are pricing a security with

a finite-horizon payoff at time T , and hence we need to encode the time remaining to maturity of the security.

Since the risk-free bond and the stock generate a complete market, every other financial security is **redundant**, meaning that its payoff can be achieved by a suitable dynamic portfolio strategy in the bond and stock. It follows from absence of arbitrage, discussed in Section 1.6, that the price $g(Q_t, t)$ must be equal to the value of the portfolio needed to replicate the same terminal payoff $G(Q_T)$. If it were not, then a strategy that would purchase the cheaper asset or portfolio while selling the more expensive one would generate immediate positive payoff without any future financial consequences.

5.3.3 Replication argument

We now develop the replication argument. In order to do that, we need to determine the portfolio positions that generate the replicating portfolio. One of the fundamental ideas of Black and Scholes (actually pointed out to them by Robert Merton, see Footnote 3 in [Black and Scholes \(1973\)](#)) that provides an answer to this problem is the concept of **dynamic hedging**. The idea is to find a particular combination of the bond and stock such that the infinitesimal return is the same as the infinitesimal return on the derivative security. Extending the infinitesimal argument to finite horizons yields the desired answer.

We develop the idea in an equivalent way, from a slightly different angle. Specifically, we want to construct a portfolio consisting of a particular combination of the stock and the derivative that makes the return on this portfolio risk-free, over an infinitesimal horizon. Since the portfolio is risk-free, it must earn the risk-free rate r , otherwise an arbitrage opportunity would emerge.

Let such a self-financing portfolio consist of one option with current price $g(Q_t, t)$ and a position of θ_t^r units of the risky stock with price Q_t . The value of this portfolio is

$$1 \cdot g(Q_t, t) + \theta_t^r Q_t.$$

By the self-financing assumption, the financial gain on the stock portion of this portfolio is $\theta_t^r dQ_t$, and the financial gain on the option portion is $1 \cdot dg(Q_t, t)$. An application of Itô's lemma implies that

$$\begin{aligned} dg(Q_t, t) &= g_Q(Q_t, t) dQ_t + \frac{1}{2} g_{QQ}(Q_t, t) (dQ_t)^2 + g_t(Q_t, t) dt \\ &= \left[g_Q(Q_t, t) \mu Q_t + \frac{1}{2} g_{QQ}(Q_t, t) \sigma^2 + g_t(Q_t, t) \right] dt + g_Q(Q_t, t) \sigma Q_t dW_t. \end{aligned}$$

The evolution of the value of the portfolio is therefore given by

$$\begin{aligned} dg(Q_t, t) + \theta_t^r dQ_t &= \left[(g_Q(Q_t, t) + \theta_t^r) \mu Q_t + \frac{1}{2} g_{QQ}(Q_t, t) \sigma^2 Q_t^2 + g_t(Q_t, t) \right] dt \\ &\quad + [g_Q(Q_t, t) + \theta_t^r] \sigma Q_t dW_t. \end{aligned}$$

We want to choose θ_t^r to make the gain on the portfolio locally risk-free, corresponding to

a zero risk exposure. This implies we must choose

$$\theta_t^r = -g_Q(Q_t, t).$$

With this choice, the financial gain on the portfolio is equal to

$$dg(Q_t, t) - g_Q(Q_t, t) dQ_t = \left[\frac{1}{2} g_{QQ}(Q_t, t) \sigma^2 Q_t^2 + g_t(Q_t, t) \right] dt. \quad (5.31)$$

At the same time, because this portfolio is risk-free, the absence of arbitrage argument implies that this portfolio then must earn the risk-free rate r , and hence

$$dg(Q_t, t) - g_Q(Q_t, t) dQ_t = r [g(Q_t, t) - g_Q(Q_t, t) Q_t] dt. \quad (5.32)$$

Equalizing the drift terms on the right-hand sides of (5.31) and (5.32), and writing Q instead of Q_t , we obtain the equation

$$rg(Q, t) = g_t(Q, t) + g_Q(Q, t) rQ + \frac{1}{2} g_{QQ}(Q, t) \sigma^2 Q^2. \quad (5.33)$$

This is a second-order partial differential equation for the price of the derivative security $g(Q, t)$. This second-order PDE has a terminal boundary condition $g(Q, T) = G(Q)$ which states that the price of the derivative security at maturity time T is equal to the payoff $G(Q)$.

An inspection of the PDE (5.33) reveals that it does not depend on the expected return on the stock μ , and hence we only need to determine parameters r and σ . This is a critical observation shown by [Black and Scholes \(1973\)](#) that simplified valuation of derivative securities dramatically. While the risk-free rate r is directly observable and the volatility of risky asset returns can be reasonably inferred from high-frequency data, measuring the expected return on a risky asset μ is an inherently difficult task.

It turns out that the replication argument combined with absence of arbitrage imply that the valuation of the derivative security does not depend on μ . This argument does not depend on the extremely simple structure of the [Black and Scholes \(1973\)](#) model and carries over to a variety of extensions as well.

5.3.4 Analytical solution

The prices of the European call and put options with payoffs (5.29)–(5.30) can be determined as closed-form expressions which only depend on quantiles of the normal distribution. These well-known formulas are provided in the following proposition.

Proposition 5.4. *Time- t prices of European call and put options with payoffs (5.29) and (5.30), respectively, with strike price K and maturity T , are given by*

$$\begin{aligned} C(Q, t) &= QN(z_1) - \exp(-r(T-t))KN(z_2) \\ P(Q, t) &= \exp(-r(T-t))KN(-z_2) - QN(-z_1) \end{aligned}$$

where $N(\cdot)$ is the cumulative standard normal distribution function, and

$$\begin{aligned} z_1 &= \frac{\log\left(\frac{Q}{K}\right) + \left(r + \frac{1}{2}\sigma^2\right)(T-t)}{\sigma\sqrt{T-t}} \\ z_2 &= z_1 - \sigma\sqrt{T-t}. \end{aligned}$$

It can be verified that $C(Q, t)$ and $P(Q, t)$ satisfy the partial differential equation (5.33) with the alternative terminal boundary conditions given by the payoffs (5.29) and (5.30), respectively.

It turns out that given a strike price K , it is sufficient to compute only the price of one of the options because the call and put option price are related through the so-called put-call parity

$$P(Q_t, t) + Q_t = C(Q_t, t) + K \exp(-r(T-t)). \quad (5.34)$$

The put-call parity result is based on a replication argument. The left-hand side of (5.34) is the value of a portfolio consisting of a put option and the stock, with payoff $\max(K - Q_T, 0) + Q_T = \max(K, Q_T)$. The right-hand side is the value of a portfolio invested in a call option and a risk-free investment with face value K , with total payoff $\max(Q_T - K, 0) + K = \max(Q_T, K)$.

Since the portfolios on both sides of the equation have identical payoffs at time T , then, by the no-arbitrage argument, they must also have the same time- t valuation. $P(Q_t, t)$ and $C(Q_t, t)$ are the prices of the options, Q_t is the stock price and $K \exp(-r(T-t))$ is the time- t value of the risk-free investment.

5.3.5 Comparison with Feynman–Kac formula and risk-neutral valuation

The Black–Scholes problem for the valuation of the derivative security yielded the partial differential equation

$$r g(Q, t) = g_t(Q, t) + g_Q(Q, t) r Q + \frac{1}{2} g_{QQ}(Q, t) \sigma^2 Q^2$$

with terminal condition $g(Q, T) = G(T)$. This equation is a special case of the PDE featured in the Feynman–Kac formula derived in Section 5.2.5. In particular, the Feynman–Kac formula implies that the solution to the Black–Scholes problem can be equivalently written as

$$g(Q_t, t) = E \left[e^{-r(T-t)} G(Q_T) \mid Q_t \right] \quad (5.35)$$

with Q_t following the dynamics

$$dQ_t = r Q_t dt + \sigma Q_t dW_t^*. \quad (5.36)$$

Equation (5.35) can be interpreted as a present discounted value formula for the time- t price of the cash flow $G(Q_T)$, discounted by a hypothetical stochastic discount factor

$$\frac{S_T^*}{S_t^*} = e^{-r(T-t)}$$

under the assumption that the dynamics of the stock price is modified to (5.36), with expected return on the stock equal to r instead of μ . We also use the notation W_t^* for the underlying Brownian motion because strictly speaking the process in (5.36) is not the same as the driving the data-generating dynamics of the stock price.

Formulas (5.35)–(5.36) reflect a result known as valuation under the risk-neutral measure. This central result in asset pricing theory states that we can view asset prices as expected cash flows discounted by risk-free rates as in (5.35) if we properly adjust the dynamics of those cash flows, as in (5.36).

Formulas (5.35)–(5.36) also provide a direct way of obtaining the analytical solution in Proposition 5.4. The price is simply the expectation of a piecewise linear function of Q_T , where Q_T is lognormally distributed. The calculation of the price therefore involves integrals over the (log)normal distribution, with quantiles determined by the kink in the payoff function.

5.3.6 Extensions and further remarks

The original Black–Scholes model was formulated for a constant interest rate and a stock price modeled as a geometric Brownian motion. One natural test of the model is to ask how well the model performs at a particular point in time in the cross-section of traded options. Since the risk-free rate r is observable for the given maturity T , the only remaining parameter that needs to be estimated is the volatility σ of the underlying stock.

One can therefore take the observed prices $C_K(Q, t)$ for options with different strike prices K , and invert the relationship to compute the so-called **implied volatility** of the stock, i.e., the volatility that is needed in the Black and Scholes (1973) model to make the price of the option consistent with the data. If the model is correct, then all options should share the same implied volatility.

This is, however, not supported by the data, and options with strike prices further away from the current stock price Q_t typically yield higher implied volatilities. Since option prices are increasing in the stock volatility due to the convex nature of their payoffs, this is equivalent to saying that the prices of such options are higher compared to what the Black–Scholes model would predict if volatility was calibrated to price correctly the options with strike prices K close to Q_t .

The literature on derivative pricing then quickly expanded to deal with these shortcomings of the seminal model, with extensions including price processes with jumps or stochastic volatility, see, for example, Cox and Ross (1976), Ross (1978), Merton (1977), or Heston (1993).

An interesting aspect of the discussion concerns the concept of market completeness

in an environment where options are traded. Consider an economy where the stock price follows an Itô process

$$dQ_t = \mu(Q_t) dt + \sigma(Q_t) dW_t$$

with a univariate Brownian motion W . The Black–Scholes model is a special case of this economy. We know that markets are complete and dynamic trading over $\mathcal{T} = [0, T]$ in the stock and the bond can replicate all square integrable time- T payoff. [Breedon and Litzenberger \(1978\)](#) present an interesting result that shows any such time- T payoff can also be replicated by time-0 static trade in call and put options with different strike prices K .

Now consider an extension where the dynamics of Q_t can also depend on a volatility factor X_t , with dynamics

$$\begin{aligned} dQ_t &= \mu(Q_t) dt + \sigma(Q_t, X_t) dW_t \\ dX_t &= \mu_X(X_t) dt + \sigma_X(X_t) dW_t \end{aligned}$$

where W_t is now a two-dimensional Brownian motion. In order to replicate complete markets, we need dynamic trade in two risky securities with linearly independent exposure vectors $\sigma(Q_t, X_t)$ and $\sigma_X(X_t)$, for example, a traded stock index and the volatility index. Typically, stock options with payoffs $G(Q_T)$ and volatility options with payoffs $G(X_T)$ are traded as well, with different strike prices. Time-0 trade in a combination of such sets of options however does not complete markets, we would need assets that allow to replicate the space of payoffs which are functions of the joint state (Q_T, X_T) , not just Q_T and X_T separately.

5.3.7 Numerical implementation

While the Black–Scholes model permits an analytical solution, extensions of the model may require numerical implementation. Next, we describe two approaches applicable in this situation. In Section 5.4, we study discretization schemes for partial differential equations of the type derived in (5.33), or, more generally, for those implied by the Feynman–Kac formula in Section 5.2.5. The discretization schemes replace derivatives for the function of interest with differences on a predetermined grid, and lead to systems of linear equations for which efficient solution methods exist, at least in cases when the state space is low-dimensional.

As a second method, we study the construction of binomial trees introduced in [Cox et al. \(1979\)](#). This is a computationally efficient method that replaces the distribution of paths of an underlying Itô diffusion using a recombining tree where each node has two alternative branches, with probabilities of reaching alternative branches that approximate the dynamics of the Itô diffusion. Evolution along the branches of the tree represents partial sums of scaled Bernoulli random variables which approximate the underlying Itô diffusion by a generalization of a result known as Donsker’s invariance principle.

5.4 Difference schemes for parabolic PDEs

In this section, we study numerical solutions of a class of partial differential equations associated with the computation of the present value of a cash flow

$$v(x, t) = E \left[\int_t^T \phi(t, s) h(X_s, s) ds + \phi(t, T) H(X_T, T) \mid X_t = x \right]$$

where

$$\begin{aligned} \phi(t, s) &= \exp \left(- \int_t^s r(X_\tau, \tau) d\tau \right) \\ dX_t &= \mu(X_t, t) dt + \sigma(X_t, t) dW_t \end{aligned}$$

where $X_t \in \mathcal{X} = (l, r)$ with an appropriately specified boundary condition at $\{l, r\}$. By the Feynman–Kac formula from Section 5.2.5, $v(x, t)$ solves the differential equation

$$-v_t(x, t) = \mathcal{D}v(x, t) + h(x, t) \quad (5.37)$$

where

$$\mathcal{D}v(x, t) = -v(x, t)r(x, t) + v_x(x, t)\mu(x, t) + \frac{1}{2}v_{xx}(x, t)\sigma(x, t)^2 \quad (5.38)$$

is a differential operator that collects derivatives with respect to the state dimension, and $v(x, T) = H(x, T)$ is the associated terminal condition. We also impose a general boundary condition in the space dimension

$$\alpha(x, t)v_x(x, t) + \beta(x, t)v(x, t) = \gamma(x, t) \quad x \in \{l, r\}, t \in [0, T] \quad (5.39)$$

for some (predetermined) functions α , β and γ that incorporates a variety of possible behaviors of X at the boundaries. We treat here the case where X is univariate but extensions to multivariate cases are straightforward.

Equation (5.37) is a PDE of the so-called parabolic type.¹ The form of these equations, together with the given terminal condition, invite a particular solution technique that resembles backward iteration on a value function.

The idea is to overlay a grid of nodes over the region $\mathcal{X} \times \mathcal{T} = [l, r] \times [0, T]$ and approximate equation (5.37) by replacing derivatives at the nodes with differences. This turns the functional equation (5.37) into an algebraic system of linear equations. This system is very easy to solve in principle but there are important considerations regarding the choice of the grid and construction of the derivatives which affect the accuracy and stability of the solution method.

¹The classification of second-order PDEs is based on the combination of coefficients $Av_{xx} + Bv_{xt} + Cv_{xt} + \text{lower order terms} = 0$. In particular, equations satisfying $B^2 - AC = 0$ are classified as parabolic, $B^2 - AC < 0$ as elliptic and $B^2 - AC > 0$ as hyperbolic. Solutions to each of the types of equations have different qualitative properties. Parabolic equations describe phenomena such as heat conduction or particle diffusion over time. Elliptic equations are used to characterize stable steady states because a local perturbation does not spread out throughout the solution. Hyperbolic equations capture finite-speed, wave-like propagation of initial perturbations.

5.4.1 Mapping of the problem to the Black–Scholes model

The Black–Scholes problem for the valuation of the derivative security is characterized by the partial differential equation

$$rg(Q, t) = g_t(Q, t) + g_Q(Q, t)rQ + \frac{1}{2}g_{QQ}(Q, t)\sigma^2Q^2 \quad (5.40)$$

with terminal condition $g(Q, T) = G(T)$. Hence we have $r(x, t) = r$, $\mu(x, t) = rx$, $\sigma(x, t) = \sigma x$, $h(x, t) = 0$, and $H(x, T) = G(x)$.

A more delicate issue is the choice of the boundary conditions at $x \in \{l, r\}$. The state space for the valuation problem is the open interval $Q \in (0, \infty)$, and does not prescribe any specific boundary conditions, since the boundaries $\{0, \infty\}$ are not part of the state space, they are so-called open.

In the implementation of the Black–Scholes model, we will not go into the details of exact characterization of these open boundaries, and instead take a practical approach. Imagine that we are interested in the solution on the interval $[a, b] \subset (0, \infty)$. We then choose boundaries $\{l, r\}$ with $0 < l < a < b < r < \infty$ such that the solution on $[a, b]$ is not sensitive to the choice of the boundaries.

In the case of option pricing with payoffs (5.29) or (5.30), a plausible observation is that if the current state Q is sufficiently distant from the strike price K at which the nonlinearity of the payoff occurs, then the price $g(Q, t)$ becomes close to linear in Q . Then we could neglect the second-order derivative term in (5.40) at those boundaries. Neglecting for simplicity also the partial derivative in time dimension $g_t(Q, t)$, we obtain the boundary condition

$$rg(Q, t) = g_Q(Q, t)rQ \quad Q \in \{l, r\}. \quad (5.41)$$

This maps to the boundary condition (5.39) with $\beta(x, t) = r$, $\alpha(x, t) = -rx$, and $\gamma(x, t) = 0$. Incorporating the time derivative $g_t(Q, t)$ back into the boundary is possible, adjusting the formulas in the following text appropriately.

Another choice to be consider is a transformation of the state variable Q . In the following, we consider the construction of an equidistant grid. At the same time, the state variable in the **Black and Scholes (1973)** model follows a geometric Brownian motion, so perhaps a grid with node distances that scale proportionally with Q may be more desirable. Alternatively, we could consider rewriting the model using $q = \log Q$ as the state variable.

To achieve that, define the transformation $f(q, t) = f(\log Q, t) = g(Q, t) = g(\exp(q), t)$. Then the derivatives are given as follows

$$\begin{aligned} f_q(q, t) &= \frac{d}{dq}g(\exp(q), t) = g_Q(\exp(q), t)\exp(q) = g_Q(Q, t)Q \\ f_{qq}(q, t) &= \frac{d}{dq}(g_Q(\exp(q), t)\exp(q)) = g_{QQ}(Q, t)Q^2 + g_Q(Q, t)Q. \end{aligned}$$




Figure 5.2: Equidistant grid on $[l, r] \times [0, T]$.

so that the partial differential equation is transformed to

$$rf(q, t) = f_t(q, t) + \left(r - \frac{1}{2}\sigma^2\right) f_q(q, t) + \frac{1}{2}\sigma^2 f_{qq}(q, t)$$

with the terminal boundary condition $f(q, t) = G(\exp(q))$. The auxiliary approximate boundary condition (5.41) can then be written as

$$f(q, t) = f_q(q, t) \quad q \in \{l, r\}.$$

We then obtain the general specification (5.37)–(5.39) with $r(x, t) = r$, $\mu(x, t) = r - \frac{1}{2}\sigma^2$, $\sigma(x, t) = \sigma$, $h(x, t) = 0$, and coefficients in the boundary condition $\alpha(x, t) = 1$, $\beta(x, t) = -1$, $\gamma(x, t) = 0$.

5.4.2 Grid construction and approximation with differences

We discretize the space and time dimensions by choosing integers I, J that define the space step $\Delta x \doteq (r - l) / I$ and the time step $\Delta t \doteq T / J$. In the space dimension, we construct a grid $\mathbb{X} = \{l, l + \Delta x, l + 2\Delta x, \dots, r - \Delta x, r\}$. Similarly, the grid in the time dimension is $\mathbb{T} = \{0, \Delta t, 2\Delta t, \dots, T - \Delta t, T\}$. Combining the two grids, we obtain a two-dimensional $(I + 1) \times (J + 1)$ node grid $\mathbb{X} \times \mathbb{T}$ on $\mathcal{X} \times \mathcal{T}$ where node (i, j) represents the state $(x, t) = (i\Delta x, j\Delta t)$. We denote values of functions on the grid as $f_{i,j} \doteq f(i\Delta x, j\Delta t)$.

Figure 5.2 depicts the construction of the grid together with the imposed terminal and boundary conditions. Nodes $v_{i,j}$, $i \in \{0, \dots, I\}$ are given by the terminal condition $v_{i,J} = H_{i,J}$. Nodes at the space boundaries $v_{0,j}$ and $v_{I,j}$, $j = \{0, \dots, J\}$ are given by the boundary condition (5.39). The goal is to approximate (5.37) at the interior nodes (i, j) .

We start with the discrete approximation of the space derivatives. We approximate the

derivative v_x at node (i, j) using values of v at adjacent nodes. We have several options:

$$\begin{aligned} \text{forward difference:} \quad v_x(i\Delta x, j\Delta t) &\approx v_{i,j}^{\bar{x}} \doteq \frac{1}{\Delta x} (v_{i+1,j} - v_{i,j}) \\ \text{central difference:} \quad v_x(i\Delta x, j\Delta t) &\approx v_{i,j}^{x_c} \doteq \frac{1}{2\Delta x} (v_{i+1,j} - v_{i-1,j}) \\ \text{backward difference:} \quad v_x(i\Delta x, j\Delta t) &\approx v_{i,j}^x \doteq \frac{1}{\Delta x} (v_{i,j} - v_{i-1,j}). \end{aligned}$$

For example, the forward difference gets its name from the fact that we are using the value of v in the 'ahead' node $(i+1, j)$. We discuss the suitability of these options later, and formalize their accuracy in Section 5.4.6.

To approximate v_{xx} we use the difference of differences. For example, we can use the backward difference of forward differences $v^{\bar{x}}$:

$$v_{xx}(i\Delta x, j\Delta t) \approx v_{i,j}^{\bar{x}\bar{x}} = \frac{1}{\Delta x} (v_{i,j}^{\bar{x}} - v_{i-1,j}^{\bar{x}}) = \frac{1}{(\Delta x)^2} (v_{i+1,j} - 2v_{i,j} + v_{i-1,j}).$$

Collecting terms, we replace the differential operator

$$\mathcal{D}v(x, t) = -v(x, t)r(x, t) + v_x(x, t)\mu(x, t) + \frac{1}{2}v_{xx}(x, t)\sigma(x, t)^2$$

at $(x, t) = (i\Delta x, j\Delta t)$ with the difference operator

$$(\mathcal{D}v)_{i,j} = -v_{i,j}r_{i,j} + v_{i,j}^{\bar{x}}\mu_{i,j} + \frac{1}{2}v_{i,j}^{\bar{x}\bar{x}}\sigma_{i,j}^2$$

(here, we used forward difference $v_{i,j}^{\bar{x}}$ as an example). We proceed in the same way at the boundaries to approximate (5.39), using a forward difference at the lower boundary $x = l$ and a backward difference at the upper boundary:

$$\begin{aligned} \alpha_{0,j}v_{0,j} + \beta_{0,j}v_{0,j}^{\bar{x}} &= \gamma_{0,j} \\ \alpha_{l,j}v_{l,j} + \beta_{l,j}v_{l,j}^x &= \gamma_{l,j} \end{aligned}$$

From these equations, we can express the values $v_{0,j}$ and $v_{l,j}$ as functions of the values at the adjacent nodes

$$\begin{aligned} v_{0,j} &= \frac{\gamma_{0,j} - (\beta_{0,j}/\Delta x)v_{1,j}}{\alpha_{0,j} - \beta_{0,j}/\Delta x} \\ v_{l,j} &= \frac{\gamma_{l,j} + (\beta_{l,j}/\Delta x)v_{l-1,j}}{\alpha_{l,j} + \beta_{l,j}/\Delta x} \end{aligned} \tag{5.42}$$

In the time dimension, we proceed in the same way, and approximate v_t using differ-



Figure 5.3: Explicit and implicit schemes for the solution of parabolic PDEs.

ences as follows:

$$\text{forward difference: } v_t(i\Delta x, j\Delta t) \approx v_{i,j}^{\bar{t}} \doteq \frac{1}{\Delta t} (v_{i,j+1} - v_{i,j}) \quad (5.43)$$

$$\text{backward difference: } v_t(i\Delta x, j\Delta t) \approx v_{i,j}^{\underline{t}} \doteq \frac{1}{\Delta t} (v_{i,j} - v_{i,j-1}) \quad (5.44)$$

The choice will determine two difference solution methods. Using the forward difference leads to the so-called implicit solution scheme, while the backward difference leads to the explicit scheme. We study each of them in turn.

5.4.3 Explicit solution scheme

The explicit scheme replaces the differential equation (5.37) with the difference scheme

$$-v_{i,j}^{\underline{t}} = (Dv)_{i,j} + h_{i,j}$$

for nodes (i, j) , $i \in \{1, \dots, I-1\}$, $j \in \{1, \dots, J\}$. We can rewrite this equation as

$$v_{i,j-1} = (\Delta t) (Dv)_{i,j} + v_{i,j} + (\Delta t) h_{i,j}. \quad (5.45)$$

The right-hand side depends only on values in nodes at time j . Hence if we know the solution $v_{\cdot,j}$ at time j at all nodes i , we can use (5.45) to compute the solution at time $j-1$. The complete solution is obtained by iterating backward on (5.45), starting with the known terminal condition $v_{\cdot,J} = H_{\cdot,J}$.

The left panel of Figure 5.3 depicts the iteration implied by the explicit scheme. Equation at node (i, j) collects information from $v_{i-1,j}$, $v_{i,j}$ and $v_{i+1,j}$ to explicitly compute the solution for $v_{i,j-1}$. For a given j , we can write the system of $I-1$ equations at nodes $i \in \{1, \dots, I-1\}$ as

$$v_{\cdot,j-1}^{\text{int}} = A_j v_{\cdot,j}^{\text{int}} + \tilde{h}_j \quad (5.46)$$

where $v_{\cdot,j}^{\text{int}} = (v_{1,j}, \dots, v_{I-1,j})'$, values $v_{0,j}$ and $v_{I,j}$ were substituted out using boundary conditions (5.42), and A_j is a tri-diagonal $(I-1) \times (I-1)$ matrix.

Stability of the explicit scheme

We want to assure that the derived iterative scheme (5.46) is stable, which means that small errors in the approximation do not propagate and amplify as we iterate backwards. We have

$$v_{:,0}^{int} = A_1 v_{:,1}^{int} + \tilde{h}_1 = A_1 A_2 v_{:,2}^{int} + A_1 \tilde{h}_2 + \tilde{h}_1 = \dots = \left(\prod_{j=1}^J A_j \right) v_{:,J}^{int} + \dots$$

This means that we want to make sure that matrices A_j have eigenvalues inside the unit circle.

For notational simplicity, we restrict our attention to the heat equation case

$$r(x, t) = \mu(x, t) = 0, \sigma(x, t) = \sigma \quad (5.47)$$

with boundary conditions

$$\alpha(x, t) = 1, \beta(x, t) = 0, \gamma(l, t) = \gamma_l, \gamma(r, t) = \gamma_r \quad (5.48)$$

In this simple case, we have

$$A_j = A = \begin{pmatrix} 1 - \frac{\Delta t}{(\Delta x)^2} \sigma^2 & \frac{\Delta t}{(\Delta x)^2} \frac{\sigma^2}{2} & 0 & 0 & \dots \\ \frac{\Delta t}{(\Delta x)^2} \frac{\sigma^2}{2} & 1 - \frac{\Delta t}{(\Delta x)^2} \sigma^2 & \frac{\Delta t}{(\Delta x)^2} \frac{\sigma^2}{2} & 0 & \dots \\ 0 & \frac{\Delta t}{(\Delta x)^2} \frac{\sigma^2}{2} & 1 - \frac{\Delta t}{(\Delta x)^2} \sigma^2 & \frac{\Delta t}{(\Delta x)^2} \frac{\sigma^2}{2} & \dots \\ 0 & 0 & \ddots & \ddots & \ddots \end{pmatrix}$$

and

$$\tilde{h}_j = \begin{pmatrix} (\Delta t) h_{i,j} + \frac{\Delta t}{(\Delta x)^2} \frac{\sigma^2}{2} \gamma_l \\ (\Delta t) h_{i,j} \\ \vdots \\ (\Delta t) h_{i,j} \\ (\Delta t) h_{i,j} + \frac{\Delta t}{(\Delta x)^2} \frac{\sigma^2}{2} \gamma_h \end{pmatrix}$$

The eigenvalues of matrix A are given by

$$\lambda_i = 1 - 2 \frac{\Delta t}{(\Delta x)^2} \sigma^2 \left(\sin \frac{i\pi}{2I} \right)^2 \quad i \in \{1, \dots, I-1\}.$$

Since we require $|\lambda_i| < 1, \forall i$, we must have

$$-1 < 1 - 2 \frac{\Delta t}{(\Delta x)^2} \sigma^2 \left(\sin \frac{i\pi}{2I} \right)^2 < 1,$$

so it is sufficient to choose Δt and Δx so that they satisfy

$$\sigma^2 \Delta t < (\Delta x)^2.$$

The explicit scheme is therefore conditionally stable. In the more general case with non-constant coefficients, it is not straightforward to characterize sufficient conditions explicitly, but the general intuition based on the stability of matrices A_j still holds. At the very least, once A_j is constructed, it is worth checking its eigenvalues to get a sense of potential stability issues.

5.4.4 Implicit solution scheme

Using the forward difference to approximate the time derivative in (5.37) leads to the difference scheme

$$-v_{i,j-1}^t = (Dv)_{i,j-1} + h_{i,j-1}$$

at nodes $(i, j-1)$, $i \in \{1, \dots, I-1\}$, $j \in \{1, \dots, J\}$, which we can write as

$$v_{i,j-1} - (\Delta t) (Dv)_{i,j-1} = v_{i,j} + (\Delta t) h_{i,j-1}.$$

This equation defines values for nodes at time $j-1$ implicitly, as an equation for three unknowns $v_{i-1,j-1}$, $v_{i,j-1}$ and $v_{i+1,j}$ collected on the left-hand side, with a known value $v_{i,j}$. The right panel in Figure 5.3 depicts the scheme. In matrix form, we now have a system of $I-1$ equations

$$A_j v_{\cdot,j-1}^{int} = v_{\cdot,j}^{int} + \tilde{h}_j$$

where we again substituted out $v_{0,j-1}$ and $v_{I,j-1}$ using the boundary conditions. This yields the iterative scheme

$$v_{\cdot,j-1}^{int} = (A_j)^{-1} \left(v_{\cdot,j}^{int} + \tilde{h}_j \right).$$

When $A_j = A$, then this scheme only requires one matrix inversion, but even when it is not, inverting the tri-diagonal matrix A_j is computationally relatively cheap.

Stability of the implicit scheme

Stability of the implicit scheme will depend on the eigenvalues of matrices $(A_j)^{-1}$. For the special case given by restrictions (5.47)–(5.48), the matrix A_j is given by

$$A_j = A = \begin{pmatrix} 1 + \frac{\Delta t}{(\Delta x)^2} \sigma^2 & -\frac{\Delta t}{(\Delta x)^2} \frac{\sigma^2}{2} & 0 & 0 & \cdots \\ -\frac{\Delta t}{(\Delta x)^2} \frac{\sigma^2}{2} & 1 + \frac{\Delta t}{(\Delta x)^2} \sigma^2 & -\frac{\Delta t}{(\Delta x)^2} \frac{\sigma^2}{2} & 0 & \cdots \\ & -\frac{\Delta t}{(\Delta x)^2} \frac{\sigma^2}{2} & 1 + \frac{\Delta t}{(\Delta x)^2} \sigma^2 & -\frac{\Delta t}{(\Delta x)^2} \frac{\sigma^2}{2} & \cdots \\ & & \ddots & \ddots & \ddots \end{pmatrix}$$

and it can be shown that $(A_j)^{-1}$ have all eigenvalues inside the unit circle. The implicit scheme for this special case is therefore unconditionally stable.

5.4.5 Upwind scheme for the choice of first difference

Another frequent stability issue concerns the choice of approximation of the first derivative. We want to study the performance of the explicit scheme, since the implicit scheme may not always be implementable when the PDE contains terms that are nonlinear in the space derivative. However, the same arguments can apply to the implicit scheme as well. In order to study this problem, we return to the PDE (5.37) and simplify it by setting

$$h(x, t) = r(x, t) = \sigma(x, t) = 0, \quad \mu(x, t) = \mu$$

This yields the first order differential equation on $\mathbb{R} \times \mathcal{T}$

$$-v_t(x, t) = \mu v_x(x, t) \quad (5.49)$$

with terminal condition $v(x, T) = H(x, T)$. As we will see, the equation describes the deterministic time propagation of a traveling wave, because the underlying Itô diffusion is now deterministic with drift μ .

We truncate the space dimension by choosing an interval $\mathcal{X} = [l, r]$ and, as before, overlay a grid $\mathbb{X} \times \mathbb{T}$ on $[l, r] \times [0, T]$ with steps $\Delta x = (r - l) / I$ and $\Delta t = T / J$. We start with the choice of forward difference to replace v_x (except at the right boundary, where we choose the backward difference).

At any interior node (i, j) , we now have the difference scheme

$$-v_{i,j}^t = \mu v_{i,j}^{\bar{x}}$$

which can be rewritten as

$$v_{i,j-1} = v_{i,j} + \mu \frac{\Delta t}{\Delta x} (v_{i+1,j} - v_{i,j}).$$

Writing this scheme in the matrix form, we obtain

$$v_{\cdot, j-1} = A^{\bar{x}} v_{\cdot, j}$$

with

$$A^{\bar{x}} = \begin{pmatrix} 1 - \mu \frac{\Delta t}{\Delta x} & \mu \frac{\Delta t}{\Delta x} & 0 & & \\ 0 & 1 - \mu \frac{\Delta t}{\Delta x} & \mu \frac{\Delta t}{\Delta x} & & \\ & & \ddots & \ddots & \\ & & & 0 & 1 - \mu \frac{\Delta t}{\Delta x} & \mu \frac{\Delta t}{\Delta x} \\ & & & 0 & -\mu \frac{\Delta t}{\Delta x} & 1 + \mu \frac{\Delta t}{\Delta x} \end{pmatrix}$$

The last line emerges from the use of backward difference at the right boundary. It is immediate to see that $A^{\bar{x}}$ has one eigenvalue equal to $1 - \mu \frac{\Delta t}{\Delta x}$ with multiplicity $I - 1$ and

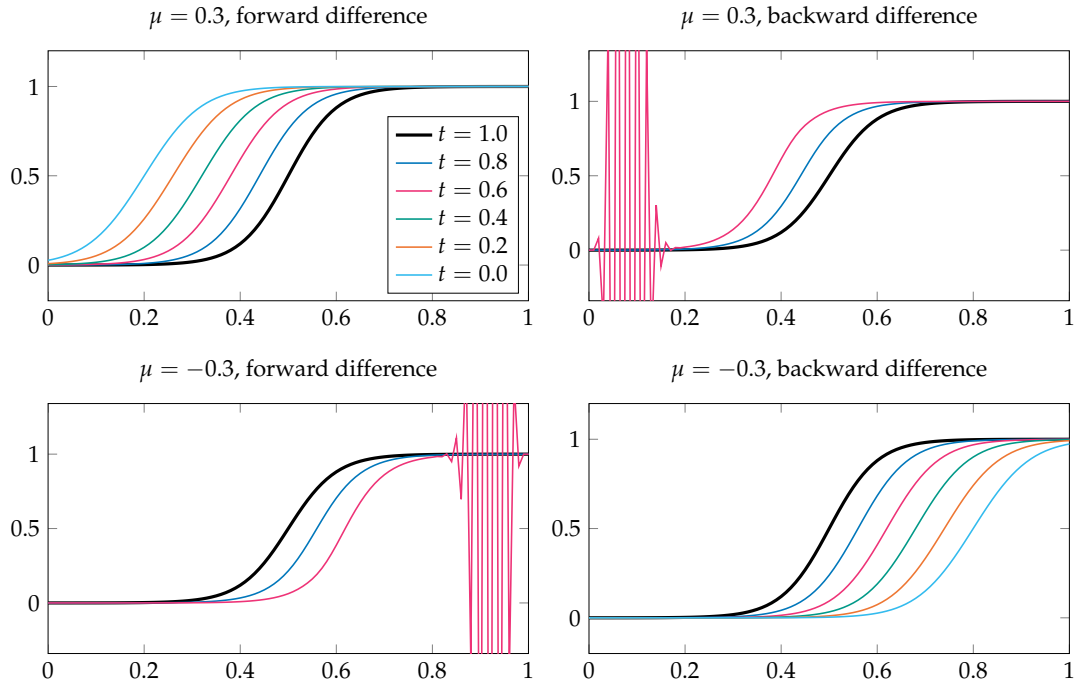


Figure 5.4: Solution to equation (5.49) with $T = 1$ and terminal condition $H(x, T) = (1 + \exp(-20(x - 0.5)))^{-1}$. The grid is chosen to be $[l, r] = [0, 1]$ with steps $\Delta x = 0.01$ and $\Delta t = 0.01$.

one eigenvalue equal to 1 with multiplicity two.

The explicit scheme with the forward difference replacing v_x will therefore be unstable when $\mu < 0$. On the other hand, when $\mu > 0$, we need $|1 - \mu \frac{\Delta t}{\Delta x}| < 1$ to achieve stability, which will be satisfied if $\Delta x > \frac{\mu}{2} \Delta t$.

The backward difference choice is a direct counterpart. Repeating the above construction, we can show that the explicit scheme with the backward difference $v^{\bar{x}}$ replacing the derivative v_x yields

$$v_{\cdot, j-1} = A^{\bar{x}} v_{\cdot, j}$$

with

$$A^{\bar{x}} = \begin{pmatrix} 1 - \mu \frac{\Delta t}{\Delta x} & \mu \frac{\Delta t}{\Delta x} & 0 & & & \\ -\mu \frac{\Delta t}{\Delta x} & 1 + \mu \frac{\Delta t}{\Delta x} & 0 & & & \\ & & \ddots & \ddots & & \\ & & & -\mu \frac{\Delta t}{\Delta x} & 1 + \mu \frac{\Delta t}{\Delta x} & 0 \\ & & & 0 & -\mu \frac{\Delta t}{\Delta x} & 1 + \mu \frac{\Delta t}{\Delta x} \end{pmatrix}.$$

The matrix $A^{\bar{x}}$ now has one eigenvalue equal to $1 + \mu \frac{\Delta t}{\Delta x}$ with multiplicity $I - 1$ and one eigenvalue equal to 1 with multiplicity two. Finally, it can also be shown that the central difference choice will always lead to instability here, regardless of the choice of μ . On the other hand, the central difference choice is always stable for the implicit scheme.

The instability can have dramatic consequences. The top row in Figure 5.4 shows the solution to the PDE with $\mu > 0$ for the case with the forward difference $v^{\bar{x}}$ on the left,

and the backward difference $v^{\bar{x}}$ on the right. The terminal condition at $T = 0$ is the thick black line, and as time t goes backward, the wave propagates to the left. While the forward difference scheme delivers an accurate solution, the solution for the backward difference scheme completely disintegrates by $t = 0.6$. The opposite situation occurs in the bottom row of Figure 5.4, for $\mu = -0.3$.

Our stability analysis revealed that it is appropriate to use the forward derivative when $\mu > 0$ (and the wave propagates to the left), and to use the backward derivative when $\mu < 0$ (and the wave propagates to the right). In computational fluid dynamics, this approach yielded the name the upwind differencing scheme because the difference at node (i, j) is informed from the direction that points against the flow of the wave.

A natural question is which difference to choose if $\mu(x, t)$ changes signs. One approach is to make the direction of the difference depend on the sign of $\mu(x, t)$. In particular, we can approximate

$$-v_t(x, t) = \mu(x, t) v_x(x, t)$$

with the difference scheme

$$-v_{i,j}^t = \mu_{i,j}^+ v_{i,j}^{\bar{x}} - \mu_{i,j}^- v_{i,j}^x$$

where $\mu_{i,j}^+ = \max(0, \mu_{i,j})$ and $\mu_{i,j}^- = \max(0, -\mu_{i,j})$.

5.4.6 Approximation errors

We now return to the discussion of approximation errors associated with replacing the derivatives in the differential equation

$$-v_t(x, t) = \mathcal{D}v(x, t) + h(x, t) \quad (5.50)$$

with differences. Establishing these approximation errors are useful because they give guidance on the how quickly approximation errors decline as we refine the grid and $\Delta x \rightarrow 0$ and $\Delta t \rightarrow 0$. Refining the grid is computationally costly, and hence it is important to understand the tradeoff between this computational cost and benefits of increasing accuracy.

Consider the second-order Taylor expansion of a function f ,

$$f(x + \Delta x) = f(x) + (\Delta x) f_x(x) + \frac{1}{2} (\Delta x)^2 f_{xx}(x) + \frac{1}{6} (\Delta x)^3 f_{xxx}(z_+)$$

where the last term represents the error of approximation, with $z \in [x, x + \Delta x]$. This yields

$$\frac{f(x + \Delta x) - f(x)}{\Delta x} = f^{\bar{x}}(x) = f_x(x) + \frac{1}{2} (\Delta x) f_{xx}(x) + \frac{1}{6} (\Delta x)^2 f_{xxx}(z) = f_x(x) + O(\Delta x)$$

and the one-sided first differences (forward and backward) approximate the derivative with accuracy $O(\Delta x)$.² As for the central difference, use the expansion of $f(x + \Delta x)$ and

²We say that a function g is 'big-O' $h(\Delta x)$ as $\Delta x \rightarrow 0$, and write $g(\Delta x) = O(h(\Delta x))$, if there exist a $\delta > 0$ and an $M > 0$ such that $|g(\Delta x)| \leq M|h(\Delta x)|$ for any $|\Delta x| < \delta$.

$f(x - \Delta x)$ to obtain

$$\begin{aligned} \frac{f(x + \Delta x) - f(x - \Delta x)}{2(\Delta x)} &= f^{x_c}(x) = f_x(x) + \frac{1}{12}(\Delta x)^2(f_{xxx}(z) + f_{xxx}(z')) \\ &= f_x(x) + O((\Delta x)^2) \end{aligned}$$

and the central difference hence approximates the derivative with accuracy $O((\Delta x)^2)$. The approximation of the first derivative using the central difference therefore achieves a higher degree of accuracy than the one-sided differences as $\Delta x \rightarrow 0$ and the grid becomes finer. However, this is only a local argument about the approximation of the derivative at a point and is not a statement about the stability of the approximation of the whole differential equation on $\mathcal{X} \times \mathcal{T}$. As we have seen in Section 5.4.5, using the central derivative may lead to global instability despite the higher degree of local accuracy.

As for the second derivative, it follows from the third-order Taylor approximation that

$$\frac{f(x + \Delta x) - 2f(x) + f(x - \Delta x)}{(\Delta x)^2} = f^{\bar{x}x}(x) = f_{xx}(x) + O((\Delta x)^2).$$

The time derivative is approximated using the backward difference for the explicit scheme and the forward difference for the implicit scheme, which implies that in both cases, the error is $O(\Delta t)$.

Combining these results, the error of approximation of the differential equation (5.50) at a particular node is of order $O(\Delta t + \Delta x)$ when we use forward or backward difference for v_x , and $O(\Delta t + (\Delta x)^2)$ when we use the central difference.³

5.4.7 Nonlinear equations and steady states

We now study a version of the neoclassical growth model

$$v(k_t, t) = \max_c E_t \left[\int_t^T e^{-\rho(s-t)} u(c_s) ds \right]$$

subject to

$$dk_t = (f(k_t) - c_t) dt + \sigma(k_t) dW_t.$$

While this is a finite horizon problem, we are mostly interested in the time-invariant solution as $T \rightarrow \infty$, which can be obtained by iterating backward in time in $t \rightarrow -\infty$. The

³The Crank–Nicholson scheme combines the implicit and explicit scheme to construct an approximation that can be interpreted as the central time difference in between the adjacent time nodes and achieves a degree of accuracy in the time dimension $O((\Delta t)^2)$. There also exist higher-order approximations of the first- and second-order derivatives in the space dimension based on the so-called Runge–Kutta methods, which utilize information from more adjacent points and achieve a higher degree of accuracy than $O(\Delta x)$ or $O((\Delta x)^2)$, respectively.

associated Hamilton–Jacobi–Bellman equation is

$$0 = \max_c u(c) - \rho v(k, t) + v_t(k, t) + v_k(k, t) (f(k) - c) + \frac{1}{2} v_{kk}(k, t) \sigma^2(k) \quad (5.51)$$

While for a given process c , this is a linear PDE, the maximization over c makes the resulting PDE for the value function nonlinear. For the choice $u(c) = (1 - \gamma)^{-1} c^{1-\gamma}$, $f(k) = Ak^\alpha - \delta k$ and $\sigma(k) = \sigma k$, we have the first-order condition

$$c(k, t) = v_k(k, t)^{-\frac{1}{\gamma}}$$

and substituting it back into (5.51), we obtain

$$-v_t(k, t) = \mathcal{D}v(k, t) \quad (5.52)$$

with

$$\begin{aligned} \mathcal{D}v(k, t) &= -\rho v(k, t) + a(k, t) v_k(k, t) + \frac{1}{2} \sigma^2 k^2 v_{kk}(k, t) \\ a(k, t) &= \frac{\gamma}{1 - \gamma} c(k, t) + (Ak^\alpha - \delta k) = \frac{\gamma}{1 - \gamma} v_k(k, t)^{-\frac{1}{\gamma}} + (Ak^\alpha - \delta k). \end{aligned}$$

Candler (2001) discusses various approaches of dealing with the nonlinearity embedded in the dependence of $a(k, t)$ on the derivative $v_k(k, t)$. Here, we provide a couple of examples.

Explicit scheme example

The explicit scheme is still conceptually easy to implement but we have to take into account stability concerns. In order to mimic the notion of upwind differencing, one proposal is $-v_{i,j}^t = \mathcal{D}v_{i,j}$ with

$$\mathcal{D}v_{i,j} = -\rho v_{i,j} + a_{i,j}^+ v_{i,j}^{\bar{k}} - a_{i,j}^- v_{i,j}^k + \frac{1}{2} \sigma^2 k_i^2 v_{i,j}^{\bar{k}\bar{k}}. \quad (5.53)$$

However, we still need to deal choose how to construct the approximation of v_k in $a(k, t)$. The insight from the linear equation is that the instability should be largely addressed by the choice $a_{i,j}^+ v_{i,j}^{\bar{k}} - a_{i,j}^- v_{i,j}^k$ in (5.53) and the choice of the difference inside $a(k, t)$ should not matter much, but since the whole equation is nonlinear, we have no theoretical result for that.

Implicit scheme example and time-invariant limits

Notice that the explicit scheme from the previous example can be written as

$$v_{i,j-1} = v_{i,j} + (\Delta t) \mathcal{D}v_{i,j}.$$

In the backward iteration scheme constructed to compute the time-invariant limit, $v_{i,j}$ is the approximation from the previous iteration, $(\Delta t) \mathcal{D}v_{i,j}$ is the ‘update’, and $v_{i,j-1}$ is the

new approximation. We are looking for $v = \lim_{j \rightarrow -\infty} v_{\cdot,j}$. However, the explicit scheme is only conditionally stable, which imposes an upper bound on Δt (as a function of Δx). Updates may therefore be small and convergence slow.

On the other hand, the implicit scheme is unconditionally stable, so that Δt can be chosen to be large. A large Δt will lead to an inaccurate solution along the transition path, but that is not a concern if we are only looking for the time-invariant solution $0 = Dv$.

As before, we can rewrite the implicit scheme $-v_{i,j-1}^t = Dv_{i,j-1}$ as

$$v_{i,j-1} - (\Delta t) Dv_{i,j-1} = v_{i,j}.$$

The problem now is that D is a nonlinear operator and therefore cannot be easily inverted to obtain $v_{\cdot,j-1}$ as a function of $v_{\cdot,j}$. One possibility is to propose to replace D with

$$D^*v_{i,j-1} = -\rho v_{i,j-1} + a_{i,j}^+ v_{i,j-1}^{\bar{k}} - a_{i,j}^- v_{i,j-1}^{\underline{k}} + \frac{1}{2} \sigma^2 k_i^2 v_{i,j-1}^{\bar{k}\underline{k}}$$

where $a_{i,j}^+$ and $a_{i,j}^-$ are taken from the node (i, j) and hence are known when constructing the discretization at $(i, j-1)$. Then

$$v_{i,j-1} - (\Delta t) D^*v_{i,j-1} = v_{i,j} \tag{5.54}$$

has a left-hand side that is linear in the unknown vector $v_{\cdot,j-1}$ and can be written as $A_j^* v_{\cdot,j-1} = v_{i,j}$ with an invertible matrix A_j^* . Notice that equation (5.54) is not a coherent description of the problem along the transition path, since we shifted the time indices of the derivatives inside $a_{i,j}$ but if it converges, it will recover the correct fixed point. For the method to work well, it is worth checking that $(A_j^*)^{-1}$ has eigenvalues inside the unit circle, but given previous arguments, there is a good chance this will be the case, even for large Δt that can speed up convergence.

[Candler \(2001\)](#) discusses strategies for dealing with nonlinear PDEs in more detail, including a two-dimensional state space, and also derives iterative schemes for policy function iteration.

5.5 Binomial trees

[Cox et al. \(1979\)](#) introduced a computationally efficient way how to approximate a geometric Brownian motion, based on a discretization of the Brownian motion on a tree, and convergence of a Binomial distribution to a Normal distribution. They then used it to price derivative securities when the underlying stock price process is a geometric Brownian motion, as in the [Black and Scholes \(1973\)](#) setup. We focus here on the simple case of the geometric Brownian motion, for extensions to more general Itô diffusions see, e.g., [Nelson and Ramaswamy \(1990\)](#). The approach can very efficiently incorporate valuation of derivatives with more complicated payoff structure.

Consider the securities market from Section 5.3 where the return on the risk-free bond




Figure 5.5: Discretization of a Brownian motion using a binomial tree.

portfolio and the stock price follow

$$\begin{aligned} dB_t &= rB_t dt \\ \frac{dQ_t}{Q_t} &= \mu dt + \sigma dW_t. \end{aligned}$$

with constant parameters r, μ, σ , where W is a Brownian motion under the data-generating measure P . Under the risk-neutral measure P^* , the stock price follows

$$\frac{dQ_t}{Q_t} = r dt + \sigma dW_t^*$$

where W^* is a Brownian motion under P^* . We are going to discretize the evolution of Q under P^* .

Figure 5.5 describes the logic of the discretization. The time interval $[0, T]$ is discretized into J periods of length $\Delta t = T/J$. In the discrete approximation of the stock price evolution, starting at time 0, the stock price Q_0 can go either up (with probability q), reaching value $u \cdot Q_0$ in Δt , or down (with probability $1 - q$), reaching value $d \cdot Q_0$ in Δt . This is repeated at every subsequent node—e.g., from node $u \cdot Q_0$, the stock price can go to $u^2 \cdot Q_0$ or $ud \cdot Q_0$.

The substantial computational advantage lies in the observation that an ‘up–down’ path ends in the same node as a ‘down–up’ path. The tree is therefore *recombining*—after j periods, there are only $j + 1$ nodes, so the size of the state vector that captures the distribution of the stock price grows linearly, rather than exponentially.

In order to construct the tree, we need to determine the parameters u, d and q . Observe that under the risk-neutral measure,

$$Q_T = Q_0 \exp \left(\left(r - \frac{1}{2} \sigma^2 \right) T + \sigma (W_T^* - W_0^*) \right)$$

hence $\log Q_T - \log Q_0 \sim N \left(\left(r - \frac{1}{2} \sigma^2 \right) T, \sigma^2 T \right)$ under P^* . We want to achieve that the

distribution of the stock price at time T on the grid matches that in the continuous-time model. On, the grid, the distribution of the number of 'up' steps after J steps is binomial with parameters (J, q) . Consequently, the distribution of $\log Q_T - \log Q_0$ on the grid is given by a (generalized) binomial distribution with mean and variance given by

$$\begin{aligned}\widehat{E}_0^* [\log Q_T - \log Q_0] &= J [q \log u + (1 - q) \log d] \\ \widehat{Var}_0^* [\log Q_T - \log Q_0] &= Jq(1 - q) (\log u - \log d)^2.\end{aligned}$$

By the Central Limit Theorem, as $J \rightarrow \infty$, this distribution converges to a normal distribution with the same mean and variance. Choosing

$$u = \exp(\sigma\sqrt{\Delta t}) \quad d = \frac{1}{u} \quad q = \frac{\exp(r\Delta t) - d}{u - d}$$

and noting that $J = T/\Delta t$, we obtain

$$\begin{aligned}\lim_{\Delta t \rightarrow 0} \widehat{E}_0^* [\log Q_T - \log Q_0] &= \lim_{\Delta t \rightarrow 0} \frac{T\sigma}{\sqrt{\Delta t}} \frac{2 \exp(r\Delta t) - u - d}{u - d} = T \left(r - \frac{1}{2} \sigma^2 \right) \\ \lim_{\Delta t \rightarrow 0} \widehat{Var}_0^* [\log Q_T - \log Q_0] &= \lim_{\Delta t \rightarrow 0} \frac{T}{\Delta t} q(1 - q) 4 (\log u)^2 = T\sigma^2\end{aligned}$$

The choice of a symmetric grid ($du = 1$) follows Cox et al. (1979) but other choices are possible as well. In particular, we could have chosen u , d and q to exactly match the mean and variance of $Q_{\Delta t}$ or its logarithm. However, the small Δt limit still holds with the above choice, and

$$\begin{aligned}\lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \widehat{E}_0^* \left[\frac{Q_{\Delta t} - Q_0}{Q_0} \right] &= r \\ \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \widehat{Var}_0^* \left[\frac{Q_{\Delta t}}{Q_0} \right] &= \sigma^2.\end{aligned}$$

as in the continuous-time model.

Once the tree is constructed, valuation proceeds as in Figure 5.6. For terminal nodes (i, J) , determine the payoff of the derivative security as $g_{i,J} = G(Q_{i,J})$. The price of the derivative security is then determined by iterating backward, using the risk-neutral pricing formula

$$g_{i,j-1} = \exp(-r\Delta t) \widehat{E}_{j-1}^* [g_{\cdot,j} \mid Q_{\cdot,j-1} = Q_{i,j-1}] = \exp(-r\Delta t) [qg_{i+1,j} + (1 - q)g_{i,j}].$$

The valuation scheme in the form of the binomial tree is rather flexible and can be used to price more complicated derivative securities than just those with a terminal payoff. For example, American options allow for early exercise before the maturity date T . Investors will choose to exercise early if the current payoff from exercising is larger than the value of the option if it held further. Using the binomial tree, we can set the value at terminal payoff date $g_{i,J} = \max(0, Q_{i,J} - K)$ and then take early exercise into account by using the




Figure 5.6: Valuation using a binomial tree

valuation formula

$$\begin{aligned}
 g_{i,j-1} &= \max \left(Q_{i,j-1} - K, \exp(-r\Delta t) \hat{E}_{j-1}^* [g_{\cdot,j}] \right) \\
 &= \max \left(Q_{i,j-1} - K, \exp(-r\Delta t) [qg_{i+1,j} + (1-q)g_{i,j}] \right).
 \end{aligned}$$

Chapter 6

Q-learning in decision problems

Textbook: *Reinforcement learning*: Sutton and Barto (2018). *Other learning methods in economics*: Sargent (1993), Sargent (1999), Evans and Honkapohja (2001)
QuantEcon: *Quantitative Economics with Python*, Topic 40 (worker search problem).

In this chapter, we discuss a method of learning that is substantially distinct from the Bayesian approach that we employed when constructing the Kalman filter in Section 4. A Bayesian learner starts with a prior probability distribution over unknown objects of the model, and uses available information that comes in the form of signals with well-described probabilistic structure to update the prior using Bayes' law.

Instead, we model an agent who learns 'optimal' actions by experimentation. The method combines aspects of Monte Carlo simulation methods with ideas from dynamic programming. Specifically, the agent learns the value of taking alternative actions in given states of the world by 'experiencing' the consequences of these actions in terms of utility flow and subsequent continuation values.

Q-learning is thus an instance of reinforcement learning, a class of algorithms in which actions are deemed to be good and thus more likely to be chosen if they lead to desirable consequences, while actions with undesirable consequences are corrected. The term Q-learning probably comes from the fact that the object to be learned in this method is a function denoted originally Q , and sometimes called the 'quality' function.

We first provide background and a general description of the method, and then apply it in the context of a worker in the McCall (1970) search problem which we modify to make the worker learn the optimal action in the form of a rule that accepts or rejects available wage offers. The application is based on the QuantEcon lecture

<https://python.quantecon.org/mccall-q.html>.

6.1 Monte Carlo methods

Monte Carlo methods are techniques used for approximation of distributions of random variables using simulations of pseudo-random draws from those distributions. These simulated distributions can then be used to evaluate moments and other statistics of those random variables. We encountered an application of the Monte Carlo method in Section 2.4.3 when we discussed the approximation of an expectations operator. Monte Carlo techniques can be advantageously used in particular in high-dimensional spaces where an explicit description of the distribution can be complicated and other numerical techniques too costly to be applied.

Consider a random variable y with distribution given by a cdf $F(y)$. If we want to numerically evaluate the expectation $E[y]$, we can create a large sample of independent draws $y^i, i = 1, \dots, I$, from the distribution $F(y)$, and then approximate

$$E[y] = \int y dF(y) \approx \frac{1}{I} \sum_{i=1}^I y^i.$$

Laws of large numbers state assumptions under which the sum converges to the expectation, and central limit theorems characterize the distributional properties of the sum.

This method can be naturally extended to dynamic environments. Consider a Markov stochastic process $s_t, t = 0, 1, 2, \dots$ that is stationary and ergodic, with values $s_t \in \mathcal{S}$, transition probability $P(s_{t+1}|s_t)$, and unconditional distribution $P(s_t)$. We are interested in evaluating the present value of income $y_t = y(s_t)$ with unconditional mean $E[y_t]$, conditional on an initial state $s_0 = s \in \mathcal{S}$,

$$V(s) = E \left[\sum_{t=0}^{\infty} \beta^t y(s_t) | s_0 = s \right].$$

One challenge when evaluating the expectation in the value function is the infinite sum. We could, for example, proceed as follows. Pick a large T and split the sum as

$$V(s) = E \left[\sum_{t=0}^T \beta^t y(s_t) | s_0 = s \right] + E \left[\sum_{t=T+1}^{\infty} \beta^t y(s_t) | s_0 = s \right].$$

Since the process s_t is stationary and ergodic, the conditional distribution of s_t converges to the unconditional one as $t \rightarrow \infty$. This allows us to approximate the last term in the above expression with the unconditional expectation

$$E \left[\sum_{t=T+1}^{\infty} \beta^t y(s_t) | s_0 = s \right] \approx \frac{\beta^{T+1}}{1-\beta} E[y(s_t)].$$

We can then proceed as follows. First, draw a large sample of paths $s_t^i, t = 0, 1, 2, \dots, i = 1, \dots, I$ initiated at $s_0 = s$, using transition probabilities $P(s_{t+1}|s_t)$ for drawing $s_{t+1}^i \sim P(s_{t+1}|s_t^i)$ to form each path i . Second, generate a large number of draws $s^j, j = 1, \dots, J$,

from the unconditional distribution $P(s_t)$. Then we can approximate

$$V(s) \approx \sum_{i=1}^I \sum_{t=0}^T \beta^t y(s_t^i) + \frac{\beta^{T+1}}{1-\beta} \frac{1}{J} \sum_{j=1}^J y(s^j).$$

6.2 Q-learning

We now combine the ideas underlying Monte Carlo simulation methods with dynamic programming concepts to recursify the calculations. For now, assume that the state space \mathcal{S} is finite.

The value function can be represented recursively using the Bellman equation

$$V(s) = y(s) + \beta E[V(s_1) | s_0 = s]. \quad (6.1)$$

We encountered this recursive representation in Chapter 1 when we studied recursive valuation of cash flows, and in Chapter 2 when we formulated the recursive problem of the McCall (1970) worker.

In both cases, we approached to problem using backward iterations. This procedure requires the evaluation of the expectation on the right-hand side of the Bellman equation (6.1) in every step of the iteration.

This expectation can be approximated using Monte Carlo draws. Generate a large number of draws s^i , $i = 1, \dots, I$, from the conditional distribution $P(s_1 | s_0 = s)$, and construct

$$y(s) + \beta E[V(s_1) | s_0 = s] \approx y(s) + \beta \frac{1}{I} \sum_{i=1}^I V(s^i). \quad (6.2)$$

Then the backward iteration algorithm can proceed as usual. Since the draws are generated from $P(s_1 | s_0 = s)$, the right-hand side of (6.2) is an unbiased estimate of the left-hand side:

$$y(s) + \beta E[V(s_1) | s_0 = s] = y(s) + \beta E \left[\frac{1}{I} \sum_{i=1}^I V(s^i) | s_0 = s \right].$$

In general, drawing a larger number of samples s^i increases the accuracy of the approximation. Instead of doing that, we now consider a conceptually different idea, and choose $I = 1$. We therefore generate a single draw of the next-period state s_1 , denoted s' . With this draw, we form the so-called temporal difference

$$TD(s, s') = y(s) + \beta V(s') - V(s). \quad (6.3)$$

This temporal difference approximates the difference between the right- and left-hand side of the Bellman equation (6.1). Since $y(s) + \beta V(s')$ is an unbiased estimate of the right-hand side of (6.1), we have

$$E[TD(s, s') | s] = y(s) + \beta E[V(s') | s] - V(s) = 0.$$

With these observations, we devise the following algorithm. Given a conjectured value function $V^n(s)$ and current state s , draw s' from $P(s'|s)$. Then construct the next iteration of the value function as $V^{n+1}(\tilde{s})$ as

$$\begin{aligned} V^{n+1}(s) &= V^n(s) + \alpha \underbrace{[y_0 + \beta V^n(s') - V^n(s)]}_{TD^n(s, s')} \\ V^{n+1}(\tilde{s}) &= V^n(\tilde{s}), \quad \tilde{s} \neq s. \end{aligned} \quad (6.4)$$

The parameter $\alpha \in (0, 1)$ is the learning rate. Once updated, repeat the iterations with s' as the new state.

The intuition behind the algorithm is as follows. If $V^n(s)$ is too low relative to the correct value, then $TD(s, s')$ will be on average positive, and the value will be on average updated upward, $V^{n+1}(s) > V^n(s)$. When $V^n(s)$ is too high, then, correspondingly, the value will be updated downward on average. In this way, the iterations will tend, under appropriate conditions, toward the correct value function.

If $V^n = V$, then $E[TD(s, s') | s] = 0$, and the value $V^n(s)$ will on average not change. However, due to the stochastic nature of the iterations, the sequence V^n never converges to a deterministic function. To assure deterministic convergence, the learning rate α has to converge to zero as $n \rightarrow \infty$ at an appropriate rate, a question we return to later.

6.2.1 Adding decision-making

The simulation algorithm for the approximation of the value function has not involved any active choice. We now extend the recursive problem to include such a possibility, which yields the Q-learning algorithm.

Consider a situation where in every state $s \in \mathcal{S}$, there is a set of available actions $a \in \mathcal{A}$. Both \mathcal{S} and \mathcal{A} are finite. The set of currently available actions \mathcal{A} can depend on s but we suppress this in the notation, and the distribution of the next period state s' can also depend on the current action a .

We then define the so-called state-action value function $Q : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ as

$$Q(s, a) = y(a) + \beta E \left[\max_{a' \in \mathcal{A}} Q(s', a') \mid s, a \right]. \quad (6.5)$$

The state-action value function is tightly related to the value function $V(s)$ through

$$V(s) = \max_{a \in \mathcal{A}} Q(s, a).$$

This reveals that the recursion for the state-action value function (6.5) is equivalent to

$$V(s) = \max_{a \in \mathcal{A}} y(a) + \beta E [V(s') \mid s, a].$$

We can now combine the temporal difference algorithm with optimal choice. For a given

state s and state-action value function Q^n , draw $s'(a)$ from the distribution $P(s'|s, a)$ for all $a \in \mathcal{A}$. With these draws, define Q^{n+1} as

$$Q^{n+1}(s, a) = Q^n(s, a) + \alpha \underbrace{\left[y(a) + \beta \max_{a' \in \mathcal{A}} Q^n(s'(a), a') - Q^n(s, a) \right]}_{\text{temporal difference } TD^n(s, a, s'(a))} \quad (6.6)$$

$$Q^{n+1}(\tilde{s}, a) = Q^n(\tilde{s}, a), \quad \tilde{s} \neq s.$$

When $Q^n = Q$, then the conditional expectation of the temporal difference is again equal to zero, $E[TD^n(s, a, s'(a)) | s, a] = 0$. Using the same intuition as in (6.4), the algorithm in (6.6) will tend to stochastically approximate the state-action value function $Q(s, a)$.

The updating rule can also be written in the form

$$Q^{n+1}(s, a) = (1 - \alpha) Q^n(s, a) + \alpha \left[y(a) + \beta \max_{a' \in \mathcal{A}} Q^n(s'(a), a') \right]$$

which indicates that the updated function is constructed as a weighted average of the old iteration and the forward-simulated continuation value, taking the next-period optimal choice a' according to the current estimate Q^n .

Once the algorithm for the approximation of the state-action value function converged to some limit \bar{Q} , we can obtain optimal value and policy as

$$a^*(s) = \arg \max_{a \in \mathcal{A}} \bar{Q}(s, a) \quad V(s) = \bar{Q}(s, a^*(s)).$$

6.2.2 Interpretation

The Q-learning algorithm learns the function Q , sometimes called the quality function, via simulation. The method is an example of so-called reinforcement learning, a class of methods in which actions that lead to high realized values or rewards are 'reinforced' as good choices by increasing their current state-action values $Q(s, a)$.

Notably, while the algorithm requires draws from the transition probabilities of the underlying state, it does not involve an explicit formation of agent's beliefs. Instead, the agent is presented with draws from the conditional distribution of the state, and learns how to take optimal action conditional on the state. The distribution of the next-period state conditional on current state and action is reflected indirectly, by more likely states being drawn more frequently.

One interpretation of Q-learning is as a computational algorithm for solving rational expectations problems. We are interested in solving problem (6.5) but evaluation of the expectations operator or the maximization operator is hard. This is advantageous when the description of the probability distribution is too complex or in situations where uncertainty is not essential but actions are complex, like in strategic games. Once the problem has converged, we obtain optimal policy under the data-generating process.

The other interpretation of the algorithm is as a description of actual behavior under

reinforcement learning. This form of learning appears markedly different from forms of learning where agents use information to update beliefs either in a Bayesian or non-Bayesian way. This requires disciplining the structure of the updating of the state-action value function and the learning rate in a behaviorally plausible way. However, the crucial question is to which extent the model of reinforcement learning leads to novel testable predictions that can be compared with data and help distinguish the model from other forms of learning.

An associated question is the embedding of the individual decision problem in an equilibrium model of market interaction. While the simulation approach appears to simplify individual decision problems by allowing to sidestep explicitly forward-looking behavior, aggregation and market clearing bring many of these considerations back into play.

6.2.3 Implementation details

There is a range of other aspects that require consideration when implementing the Q-learning algorithm.

First, the algorithm in the form (6.6) requires the state space \mathcal{S} and action space \mathcal{A} to be finite. This form of the algorithm is known as tabular Q-learning, since it relies on updating individual values in the matrix $Q(s, a)$ organized in a table. If \mathcal{S} or \mathcal{A} are infinite, for example, continuous, then it must be discretized to turn the problem into the tabular form.

Alternative methods for continuous state spaces use projection methods for the functions V^n and iterate on corresponding projection coefficients. The machine learning literature sometimes utilizes neural networks for complex projection problems, and calls the algorithm deep Q-learning.

Second, the rate of convergence of the learning rate α toward zero as $n \rightarrow \infty$ involves the tradeoff between computational costs and ability to converge to the correct solution. When α vanishes to zero too slowly, computational costs will be high. On the other hand, if it vanishes too quickly, the solution may not converge to the true solution. For more detail on these convergence results, see Sutton and Barto (2018) or Li et al. (2023).

Third, algorithms of the form (6.6) applied to complex problems can get stuck in local maxima when Q^n leads to suboptimal actions that are self-confirming because optimal actions have lower state-action values and are thus not explored. This is a problem also known from Markov chain Monte-Carlo methods in econometrics. A solution to such a problem involves experimentation in which values determined by actions chosen via $\max_{a' \in \mathcal{A}} Q^n(s', a')$ are occasionally replaced by $Q^n(s', \tilde{a})$ with \tilde{a} drawn randomly. As with the learning rate, the experimentation rate must ultimately decay to zero for the algorithm to converge deterministically, and the decay rates of the experimentation rate and learning rate toward zero must be suitably balanced.

Fourth, updating will be infrequent and hence slow for parts of the state spaces that are only infrequently visited. Learning the value function in such parts of the state space may benefit from re-starting the simulations in these states.

Fifth, initialization of the Q-learning algorithm requires a suitable choice of the function $Q^0(s, a)$, especially in a situation when the algorithm is meant to describe actual behavior under reinforcement learning.

Sixth, when Q^n is not the correct state-action value function, the algorithm in (6.6) will tend to overstate the updates due to the convexity of the max operator, yielding an upward bias in the value function. To understand this, consider a Q^n for which each of the value $Q^n(s, a)$ are given by

$$Q^n(s, a) = Q(s, a) + \varepsilon(s, a)$$

where $\varepsilon(s, a)$ is zero-mean noise. Then maximization of $Q^n(s, a)$ over a will typically lead to larger values than maximization of $Q(s, a)$. As above, the convergence rate of learning must be balanced to make this bias vanish in the limit $n \rightarrow \infty$.

6.3 Q-learning in the worker search problem

We now implement the Q-learning algorithm in a simple application based on the [McCall \(1970\)](#) worker search model. The application utilizes the corresponding [Quantecon](#) lecture

<https://python.quantecon.org/mccall-q.html>.

The model environment is the same as studied in Section 2.1, and we briefly review it here.

The worker lives in an infinite-horizon economy, with discrete time $t = 0, 1, 2, \dots$. Every period t , an iid wage offer w is drawn from distribution $F(w)$, with $F(0) = 0, F(B) = 1$ for some $B > 0$. Since the baseline tabular Q-learning algorithm requires a finite state space, we will assume throughout that the support of the wage offer distribution is finite, given by a discrete grid $w^i, i = 1, \dots, I$, with probability distribution $\hat{f}^i, i = 1, \dots, I$. To keep notational consistency, we continue writing

$$E[v] = \int_0^B v(w) dF(w) = \sum_{i=1}^I v(w^i) \hat{f}^i.$$

This discrete distribution could have been constructed from an underlying continuous distribution using quadrature methods described in Section 2.4.3. Alternatively, we could rely on projection methods applied to the state-action value function and iterations based on a ‘deep Q-learning’ algorithm to update the projection coefficients.

The worker chooses to accept or reject the offer, $a_t \in \{\text{accept, reject}\}$. Acceptance means that the worker receives income $y_t = w$ forever. Rejection implies that the worker receives unemployment benefit $y_t = c$ and moves to next period where a new offer is drawn. Time is discounted at rate $\beta \in [0, 1)$.

The worker hence solves the sequence problem

$$V_0^* = \max_{\{a_t\}_{t=0}^{\infty}} E_0 \left[\sum_{t=0}^{\infty} \beta^t y_t \right] \quad (6.7)$$

where $a_t \in \{\text{accept, reject}\}$ if the worker has not yet accepted any earlier offer, and $a_t \in \{\}$ otherwise. V_0^* denotes the value function, and we assume that V_0^* conditions on the initial offer w_0 being observed.

Every decision a_t is made conditional on the time- t information set, which contains the history of all offers up to time t , $w^t = (w_0, \dots, w_t)$. The problem of a worker with current offer w at hand can be formulated recursively as

$$V(w) = \max_{\{\text{accept, reject}\}} \left\{ V^a(w), c + \beta \int_0^B V(w') dF(w') \right\}$$

where $V^a(w)$ is the value of accepting the offer.

In this formulation, we assumed that once an offer w is accepted, the worker works at that wage forever. The value of accepting an offer thus is

$$V^a(w) = \frac{w}{1 - \beta}.$$

However, as we have seen in Section 2.2.1, in this iid environment this assumption can be dropped without any consequence, and the problem modified as follows. In every period, the worker can decide whether to continue working at the same wage w or leave to unemployment, in which case a new wage offer arrives in the next period. This is equivalent to offering to the worker the same wage contract at wage w every period.

The value of accepting an offer can then be written as follows

$$V^a(w) = w + \beta \max_{\{\text{accept, reject}\}} \left\{ V^a(w), c + \beta \int_0^B V(w') dF(w') \right\}$$

Since the wage draws are iid, the outside option is constant over time. Hence, even if the worker is allowed to leave the current job that guarantees wage w which was previously accepted, such an option would never be exercised. In other words, if the wage was sufficiently high to be accepted in one period, it is sufficiently high to be accepted in any other period.

This conclusion would be different in a model with an additional persistent state variable that affects the outside option, for example in the form of a predictable distribution of future offers or their frequency. It is also worth stressing that this option to leave is distinct from a model of search on the job in which the worker samples new wage offers while continuing to work in the existing job, and only accepts an offer if it is better than the current wage.

6.3.1 Implementing Q-learning

In the McCall (1970) model, the worker understands the probabilistic structure of the model, which allows to form expectations (subjective or objective) over the next-period offers.

Instead, we now implement the Q-learning approach described earlier. Under this

approach, the worker does not have available the probability distribution of next-period offers. Instead, the worker observes realized wage offer draws and makes accept or reject decisions. The Q-learning algorithm is an example of reinforcement learning algorithms in which the worker is rewarded for making decisions that, over time, lead to high payoffs. Through this process, the worker learns the value of alternative actions in a given state, which then allows to deduce optimal action.

We start by rewriting the problem in the form of the state-action value function, defined in its general form in (6.5). In our context, the function $Q(w, a)$ is given by

$$\begin{aligned} Q(w, \text{accept}) &= w + \beta \max_{\{\text{accept}, \text{reject}\}} \{Q(w, \text{accept}), Q(w, \text{reject})\} \\ Q(w, \text{reject}) &= c + \beta \int_0^B \max_{\{\text{accept}, \text{reject}\}} \{Q(w', \text{accept}), Q(w', \text{reject})\} dF(w'). \end{aligned}$$

The equations highlight the distinction between w and w' in the next-period offers. The first line is the value of accepting an offer, which results in working for one period, and then an option to continue at the same wage w or leaving to unemployment. The second line is the value of rejecting the offer, which results in receiving the unemployment benefit c and then drawing a new wage w' from the distribution $F(w')$ at the beginning of next period.

This state-action value function maps to the value functions $V^a(w)$ and $V(w)$ as follows:

$$V^a(w) = Q(w, \text{accept}) \quad V(w) = \max_{\{\text{accept}, \text{reject}\}} \{Q(w, \text{accept}), Q(w, \text{reject})\}.$$

Let us simplify the notation for the state-action value function and write it as

$$\begin{aligned} Q(w, \text{accept}) &= w + \beta \max_{a' \in \mathcal{A}} Q(w, a') \\ Q(w, \text{reject}) &= c + \beta \int_0^B \max_{a' \in \mathcal{A}} Q(w', a') dF(w') \end{aligned}$$

for $\mathcal{A} = \{\text{accept}, \text{reject}\}$.

Recall that the problem for the state-action value function requires the evaluation of an integral. As indicated earlier, one possibility is to apply a Monte-Carlo approach, draw a large number I of offers $w^j, j = 1, \dots, I$ from the distribution $F(w')$, and then approximate

$$\int_0^B \max_{a' \in \mathcal{A}} Q(w', a') dF(w') \approx \frac{1}{I} \sum_{i=1}^I \max_{a' \in \mathcal{A}} Q(w^i, a').$$

Because the integral is an expectation over the distribution $F(w')$, the right-hand side is an unbiased estimate of the integral, for any value of I . A law of large numbers then guarantees convergence of the sum to the integral.

If we solved the problem numerically using backward iteration, we would aim at choosing at high I to get an accurate evaluation of the integral. Instead, we do the op-

positive, choose $I = 1$, and denote the single draw of the next period wage w' . Then the sum, now consisting of a single summand, is still an unbiased, albeit a very inaccurate, estimate of the integral. Then we have

$$\begin{aligned} Q(w, \text{accept}) &= w + \beta \max_{a' \in \mathcal{A}} Q(w, a') \\ Q(w, \text{reject}) &\approx c + \beta \max_{a' \in \mathcal{A}} Q(w', a') \end{aligned}$$

where the second equation holds in expectation.

Following equation (6.3), we form the temporal differences for accept and reject decisions

$$\begin{aligned} TD^n(w, \text{accept}) &= w + \beta \max_{a' \in \mathcal{A}} Q^n(w, a') - Q^n(w, \text{accept}) \\ TD^n(w, \text{reject}) &= c + \beta \max_{a' \in \mathcal{A}} Q^n(w', a') - Q^n(w, \text{reject}). \end{aligned}$$

If the candidate solution $Q^n(w, a)$ is the correct state-action value function $Q(w, a)$, then we must have $TD^n(w, \text{accept}) = 0$ and $E[TD^n(w, \text{reject})] = 0$.

The intuition underlying the construction of the temporal differences is the same as in the general case. For example, when the function $Q^n(w, \text{accept})$ for the current state and action (w, accept) on the left-hand side is too low relative to the continuation policy $w + \beta \max_{a' \in \mathcal{A}} Q^n(w, a')$ on the right-hand side. Then the temporal difference is positive, $TD^n(w, \text{accept}) > 0$. Similarly, if $TD^n(w, \text{reject})$ is too low relative to the continuation policy $c + \beta \max_{a' \in \mathcal{A}} Q^n(w', a')$, then the temporal difference $TD^n(w, \text{reject})$ is positive at least on average, $E[TD^n(w, \text{reject})] > 0$.

Choosing a learning rate parameter $\alpha \in (0, 1)$, the Q-learning algorithm updates the state-action value function in tabular form as follows:

$$\begin{aligned} Q^{n+1}(w, a) &= Q^n(w, a) + \alpha TD^n(w, a) \\ Q^{n+1}(\bar{w}, a) &= Q^n(\bar{w}, a), \quad \bar{w} \neq w. \end{aligned}$$

The updating algorithm relies on the above-described intuition. If $Q^n(w, a)$ is too low, then the updating algorithm updates the value upward, and vice versa if $Q^n(w, a)$ is too high. Hence on average, the algorithm should converge in a stochastic sense to the correct state-action value function.

6.3.2 Implementation details

An implementation of the algorithm requires taking care of additional aspects of the numerical algorithm which we mention here without providing details or proofs.

First, when the distribution of wage offers $w \in [0, B]$ has a continuous support, then for the purposes of the algorithm, it must be discretized to a finite grid with nodes w^1, \dots, w^I . The distribution $F(w)$ is then replaced with its discrete counterpart with probabilities $\hat{f}^i = f(w^i)$, $i = 1, \dots, I$. The state-action value function is then defined on the grid points

$Q(w^i, a)$. An alternative is to use a projection method to approximate the function $Q(w, a)$ using a set of basis functions, and use the Q-learning algorithm to update the projection coefficients. Methods based on this idea are sometimes called deep Q-learning methods.

Second, the temporal differences $TD^n(w, a)$ are random because they depend on the sequential draws of $w' \sim F(w')$. This means that for a fixed $\alpha \in (0, 1)$, the sequence $Q^n(w, a)$, $n = 1, 2, \dots$ only converges in a stochastic sense. To obtain a sequence that converges deterministically, we must choose $\alpha_n \rightarrow 0$ as $n \rightarrow \infty$. The rate at which α_n vanishes to zero cannot be too high, otherwise the algorithm would not have time to converge.

Third, algorithms of this type can get stuck in local maxima for which the state-action value function is such that it dictates the incorrect action and the benefits of the correct action that would lead to a higher value function then would never be explored. For this reason, it is sensible to add to the algorithm an aspect of experimentation. Specifically, in every step n choose $\varepsilon_n \in (0, 1)$ such that the temporal difference $TD^n(w, a)$ is constructed as above with probability $1 - \varepsilon_n$. With probability ε_n , the worker 'experiments' and chooses the alternative action, so that

$$\begin{aligned} TD^n(w, \text{accept}) &= w + \beta \min_{a' \in \mathcal{A}} Q^n(w, a') - Q^n(w, \text{accept}) \\ TD^n(w, \text{reject}) &= c + \beta \min_{a' \in \mathcal{A}} Q^n(w', a') - Q^n(w, \text{reject}). \end{aligned}$$

This modification allows exploration of alternative options that may be left out otherwise. Again, for the algorithm to ultimately converge, the rate of experimentation has to vanish to zero as $n \rightarrow \infty$.

We refer to the textbook treatment in [Sutton and Barto \(2018\)](#) for additional discussion of these implementation details.

6.3.3 Results

We study an implementation in a particular calibrated version of the [McCall \(1970\)](#) economy. Wage offers are drawn from the distribution $Beta(1.2, 1.2)$, discretized to $N = 30$ nodes. The value of unemployment is $c = 0.1$, and the time preference parameter is $\beta = 0.96$. We choose exponential decay rates for learning and experimentation

$$\alpha_n = \bar{\alpha}_0 e^{-n\bar{\alpha}} \quad \varepsilon_n = \bar{\varepsilon}_0 e^{-n\bar{\varepsilon}}$$

with $\bar{\alpha}_0 = 1$, $\bar{\alpha} = 10^{-4}$, $\bar{\varepsilon}_0 = 0.1$, $\bar{\varepsilon} = 10^{-3}$. A worker who accepts a wage offer is allowed to quit later.

The results for the baseline model are presented in [Figure 6.1](#). The black squares show the true value function $V(w)$ in the left panel and decision rule $a(w)$ in the right panel obtained by solving the recursive equation for the present value of the reservation wage (2.20) in [Section 2.4.1](#). The value $a(w) = 1$ corresponds to the acceptance decision. As expected, the optimal policy is characterized by a reservation wage \bar{w} above which the wage offers are accepted. The grey area in the left panel depicts the shape of the wage offer distribution.

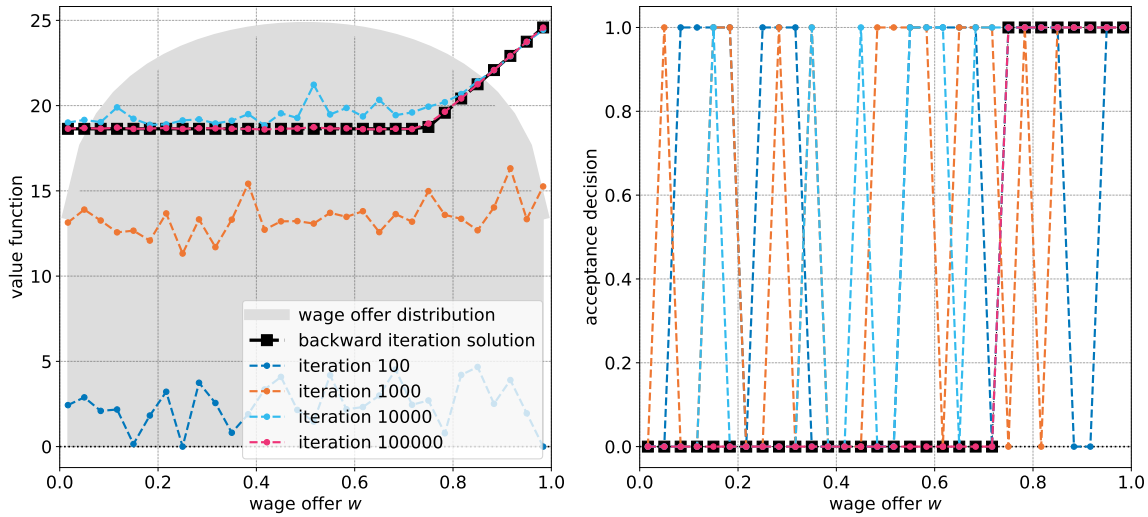


Figure 6.1: Results for the Q-learning algorithm in the worker search model for the baseline parameterization. The left panel shows the value function $V^n(w)$, the right panel the policy choice $a^n(w)$.

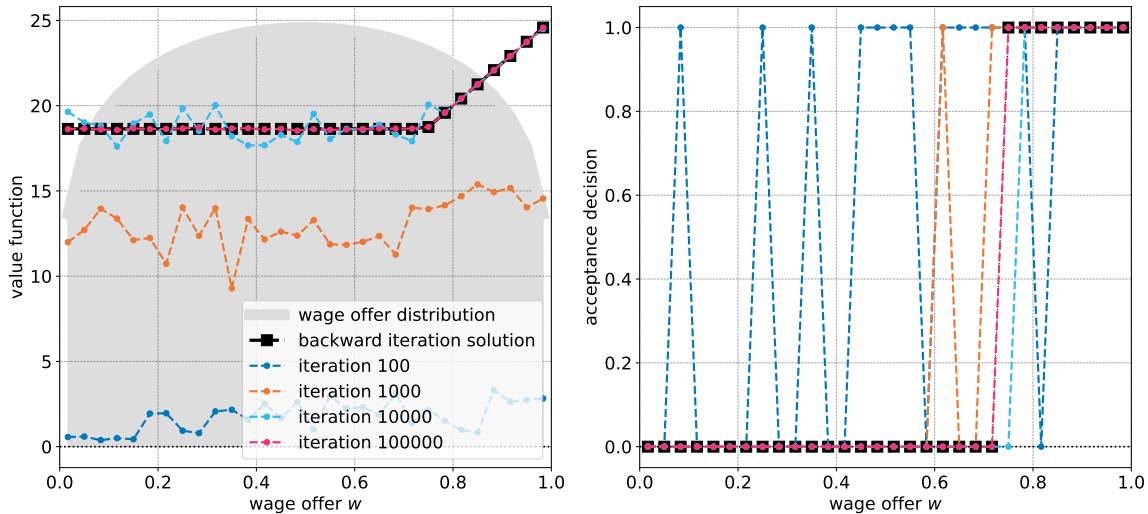


Figure 6.2: Results for the Q-learning algorithm in the worker search model for the case when leaving an accepted job is not possible.

The left panel of Figure 6.1 displays the convergence of the value function iterates $V^n(w) = \max_a Q^n(w, a)$ as n increases. After 10^4 iterations, the value function overshoots, a mechanism described in Section 6.2.3. Then, as the learning rate vanishes to zero, the value function $V^n(w)$ converges to the true value function. The decisions in the right panel also converge. However, even after 10^4 iterations, there are plenty of errors in the decision rule.

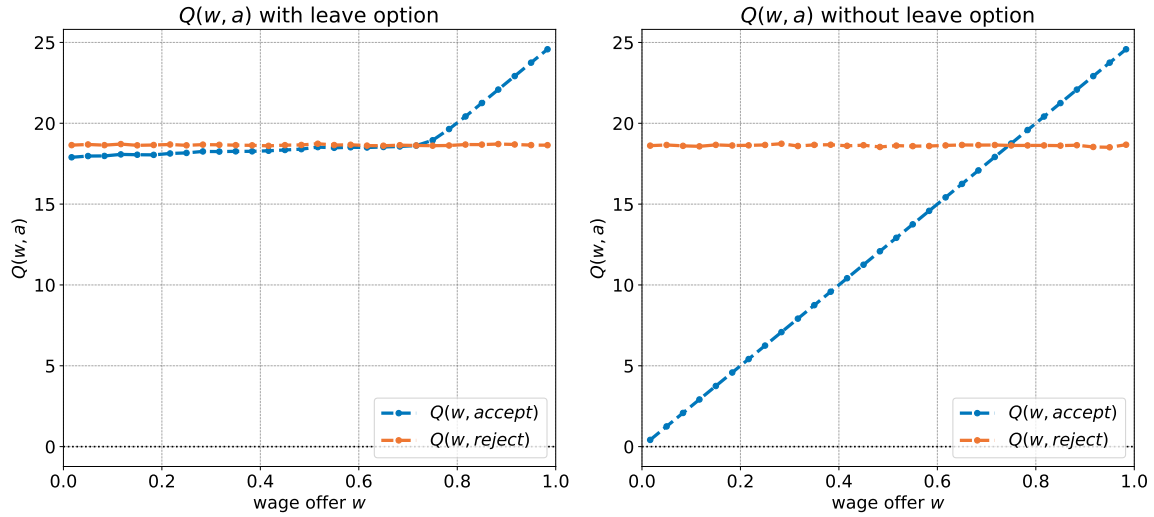


Figure 6.3: The effect of option to leave an accepted job on the values of accepting and rejecting an offer.

Figure 6.2 shows that when the option to leave an accepted job is not available, the convergence in the value function $V^n(w)$ is initially slower but without the subsequent overshooting, and the decision rules $a^n(w)$ become more accurate sooner.

The reason for the faster convergence of the decision rule under the no-quitting rule is explained in Figure 6.3. When the worker is allowed to quit, then the value of accepting the offer is

$$Q(w, \text{accept}) = w + \beta \max_{a' \in \mathcal{A}} Q(w, a')$$

while when the quit option is not available, the value of accepting is

$$Q(w, \text{accept}) = w + \beta Q(w, \text{accept}) = \frac{w}{1 - \beta}.$$

This means that with the option to quit, incorrectly accepting an offer that should have been rejected is only a one-period mistake that can be corrected next period. This means that for offers that should be rejected, i.e., those with $w < \bar{w}$, the values $Q(w, \text{accept})$ and $Q(w, \text{reject})$ are close to each other, reflecting the fact that welfare consequences of such a mistake are small. Due to the stochastic nature of the iterations, it is therefore more likely to happen that $Q^n(w, \text{accept}) > Q^n(w, \text{reject})$ even for wages $w < \bar{w}$. The differences in $Q^n(w, a)$ for $n = 10^5$ are shown in the left panel of Figure 6.3.

On the other hand, the welfare consequences of accepting a job with $w < \bar{w}$ when the option to quit is not present are large, as depicted in the right panel of Figure 6.3. Due to the large distance between the values $Q^n(w, \text{accept})$ and $Q^n(w, \text{reject})$, it is less likely that these will be ranked incorrectly in the stochastic iterations, and the decision rule is more accurate.

Figure 6.4 depicts the learning dynamics for the case when wage offers are *Beta* (5,2)

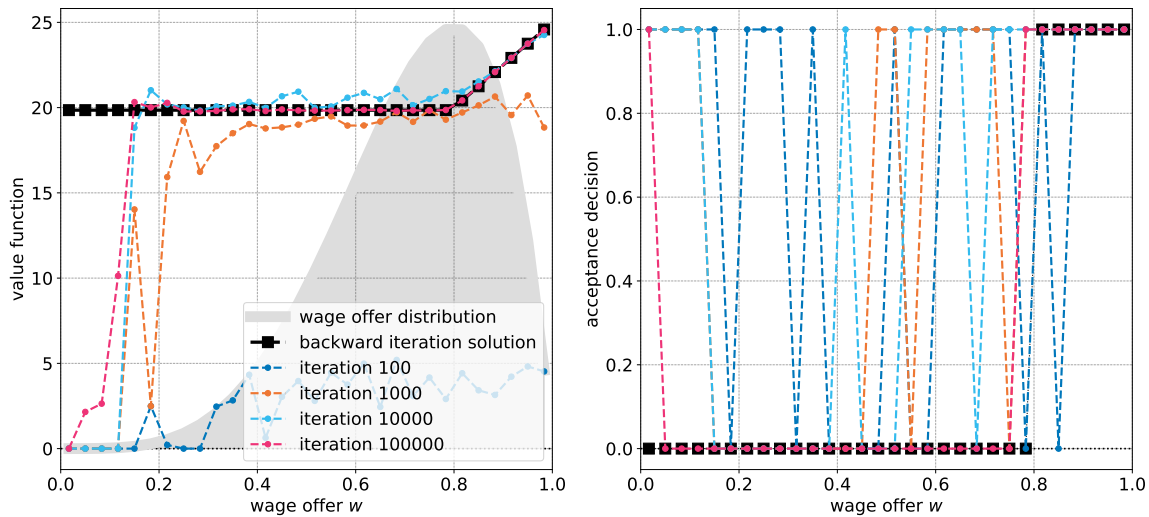


Figure 6.4: Learning the state-action value function for a skewed offer distribution.

distributed. The wage offer distribution are negatively skewed, and there is very little mass in the distribution below $w = 0.2$. Consequently, that part of the state space is visited only infrequently, and learning for those wage offers proceeds only very slowly, not reaching meaningful accuracy even after 10^6 iterations.

6.4 Summary

Q-learning and other forms of reinforcement learning is a fundamentally different approach to learning compared to learning based on Bayesian and non-Bayesian updating of beliefs. It sets aside the probabilistic structure of the problem and stipulates that agents should choose actions that had positive utility consequences in the past. High rewards reinforce choices leading to those rewards. The absence of a probabilistic model for the exogenous state dynamics also leads this algorithm to be called ‘model-free’, despite the preserved dynamic structure of the problem.

Under suitable assumptions on the model structure, decisions can converge to those obtained under rational expectations. The process of learning can be very slow, which is the consequence of the unstructured form of learning, as well as emerging from the fact that the learning process has a local character. Observing a random draw of the next-period state for the given current value of the state variable and observing the consequence of an action updates the state-action value function locally for that value of the current state or in its vicinity if a projection method is used to update the state-action value function. This is distinct from Bayesian learning in which the agent can understand the consequences of the learned information for the decision-making in other parts of the state space.

The presented Q-learning abstracts from many interesting related aspects. One category is computational, and relates to characterization of the speed of convergence of learn-

ing, choice of state space or the choice of projection method. The theoretical literature providing such results is growing quickly. The Sutton and Barto (2018) *'Reinforcement Learning'* textbook provides useful starting points to such analysis. The other is more economic, and involves embedding of the individual decision problem into models of strategic interaction and equilibrium dynamics, and studying policy implications.

Ultimately, the success of this form of learning will be deemed either in terms of its ability to approximate and solve economic models, for example with rational expectations, that are not tractable with existing methods, or in terms of its ability to deliver novel predictions for choice behavior that are distinct from the existing forms of learning, be it Bayesian or non-Bayesian.

In this respect, the existing literature on learning in economics, including non-Bayesian models, is very large. For insights into the intellectual history of non-Bayesian learning in macroeconomics, it is useful to start with Sargent (1993) *'Bounded Rationality in Macroeconomics'*, Sargent (1999) *'The Conquest of American Inflation'*, or Evans and Honkapohja (2001) *'Learning and Expectations in Macroeconomics'*. Evans et al. (2005) provide an inspiring interview.

Appendix A

Elements of probability theory

In this appendix, we provide foundations to the theory of stochastic processes that will be particularly useful in the framework of macroeconomic modeling and asset pricing. We therefore proceed by setting up a specific environment suitable for our purposes, and then fill in the theory. The central goal is to study the class of **Markov processes** which is closely related to problems solved using dynamic programming methods. Some of the concepts introduced in this section are more comprehensive than what we will subsequently need.

A.1 States, paths, partial histories and probabilities

We consider an infinite-horizon discrete-time environment with time indexed by $t \in \mathcal{T} = \{0, 1, 2, \dots\}$. Finite-horizon setups can be constructed by simply truncating the set of time indices at $T < \infty$. In every period t , the economy will find itself in **state** s_t drawn from a finite set $\mathcal{S} = \{1, 2, \dots, S\}$, i.e., $s_t \in \mathcal{S}$. The states can be completely abstract descriptions of the world like ‘rainy day’, ‘sunny day’, ‘expansion’, ‘recession’, and so on. The assumption of a finite state space is made here for simplicity but can be directly extended to infinite state spaces, at the cost of increasing the mathematical and notational complexity.

An infinite-horizon sequence of these states constitutes a **path** $s^\infty \doteq (s_0, s_1, s_2, \dots)$, i.e., a particular trajectory along which the economy evolves. We will denote the set of all such paths Ω and call it the **sample space** (and denote generic elements of this sample space $\omega \in \Omega$). The truncated sequence of states $s^t \doteq (s_0, s_1, \dots, s_t)$ constitutes a **partial history** for the economy. Observe that at time t , there are S^{t+1} such distinct partial histories that can potentially be realized.

We now want to assign probabilities to paths. Observe that there are infinitely many paths s^∞ , so typically each individual path will have probability zero of occurring. However, by assigning each individual path a zero probability, we do not get very far. Consider the following simple example. Imagine that the realization of each state every period is equally likely (i.e., $P(s_t = s) = S^{-1}, \forall s \in \mathcal{S}$) and independent over time, then each path s^∞ is also ‘equally likely’ to occur. Since we know that one path has to be realized for sure, the probabilities of each one of the paths occurring have to sum up to one. However, there

are infinitely many of them, so we are facing a “ $0 \cdot \infty$ ” problem (infinitely many paths each occurring with probability zero). Clearly, this is not well defined and this construction is therefore not useful.

Instead, we will assign probabilities to **sets** of paths. In order to do that, we need to decide which sets we can assign probabilities to. There are various ways of doing that, and certain rules need to be followed for probabilities to be consistent but one secure way in our environment is the following. Let $F(s^t)$ be the set of paths s^∞ that share the common partial history s^t , formally

$$F(s^t) = \{s^\infty = (\tilde{s}_0, \tilde{s}_1, \tilde{s}_2, \dots) \in \Omega : (\tilde{s}_0, \tilde{s}_1, \dots, \tilde{s}_t) = s^t\}.$$

Observe that for a given t , there are S^{t+1} such distinct sets (one for each distinct partial history), and these sets are disjoint. We say that s^{t+j} is a **continuation** of s^t if the first $t+1$ states (s_0, s_1, \dots, s_t) of s^{t+j} are the same as s^t .

We now define a **set of all sets of paths** to which we can assign probabilities, called a **σ -algebra** and denoted \mathcal{F} , according to the following definition:

Definition A.1. Let Ω be a given set. Then a **σ -algebra** \mathcal{F} on Ω is a family \mathcal{F} of subsets of Ω that satisfies:

1. $\emptyset \in \mathcal{F}$, i.e., the empty set belongs to \mathcal{F}
2. if a set $F \in \mathcal{F}$, then its complement $F^C \in \mathcal{F}$ (closure to complements)
3. if sets $F_1, F_2, \dots \in \mathcal{F}$, then $\bigcup_{i=1}^{\infty} F_i \in \mathcal{F}$ (closure to countable unions).

The σ -algebra \mathcal{F} that interests us is built in the following way. First include all sets of the form $F(s^t)$ for all t and all s^t , i.e., all sets of paths that share a common partial history s^t . The rest consists of applying the three rules from Definition A.1. We first include the empty set. Then we include complements of all sets that are already in \mathcal{F} , and countable unions of all sets that are already in \mathcal{F} , and repeat the application of rules 2. and 3. until we cannot add anything else. Observe that each element F of an (arbitrary) σ -algebra \mathcal{F} is a subset in Ω but not necessarily vice versa. Set $F \in \mathcal{F}$ are called **events**, while elements $\omega \in \Omega$ (in our case, these are the paths s^∞) are called **outcomes**.

We can now assign to each set $F \in \mathcal{F}$ the **unconditional probability** $P(F)$. A probability measure P is a function that takes elements F from \mathcal{F} (i.e., specific sets in Ω) and assigns to each F a number in $[0, 1]$. In order for the probability measure to be well-defined, it must satisfy certain elementary rules:

Definition A.2. A **probability measure** P on (Ω, \mathcal{F}) is a function $P : \mathcal{F} \rightarrow [0, 1]$ such that

1. $P(\emptyset) = 0, P(\Omega) = 1$
2. If $F_1, F_2, \dots \in \mathcal{F}$ and $\{F_i\}_{i=1}^{\infty}$ are disjoint (i.e., $F_i \cap F_j = \emptyset$ for $i \neq j$) then $P(\bigcup_{i=1}^{\infty} F_i) = \sum_{i=1}^{\infty} P(F_i)$.

For the sets $F(s^t)$ of all paths with common partial history s^t , we will use the shorthand notation $P(s^t)$ instead of $P(F(s^t))$. Since sets $F \in \mathcal{F}$ can be assigned probabilities (can be 'measured') we call them **measurable sets**.

We can use the **Bayes theorem** to construct **conditional probabilities** from unconditional ones. Specific examples that interest us condition on individual partial histories. In particular, let s^t be a time- t partial history that has a positive unconditional probability $P(s^t)$, and s^{t+j} a time $t+j$ partial history. Then, if s^{t+j} is a continuation of s^t , we have

$$P(F(s^{t+j})|F(s^t)) = \frac{P(F(s^{t+j}) \cap F(s^t))}{P(F(s^t))} = \frac{P(F(s^{t+j}))}{P(F(s^t))}$$

or, in shorthand notation,

$$P(s^{t+j}|s^t) = \frac{P(s^{t+j})}{P(s^t)}.$$

On the other hand, if s^{t+j} is not a continuation of s^t , we obtain

$$P(s^{t+j}|s^t) = \frac{P(F(s^{t+j}) \cap F(s^t))}{P(F(s^t))} = \frac{P(\emptyset)}{P(F(s^t))} = 0.$$

We will be particularly interested in cases when probabilities conditional on partial histories s^t are equivalent to those conditional only on current state s_t .

Remark A.1. Occasionally, we will want to represent the probability $P(s^t)$ by explicitly integrating across all paths $s^\infty \in F(s^t)$. Formally, we cannot simply sum up the probabilities of individual paths s^∞ because generically, $P(s^\infty) = 0$. We will therefore use the notation

$$P(s^t) \doteq P(F(s^t)) = \int_{\omega \in F(s^t)} dP(\omega)$$

where we properly integrate over the set of paths $F(s^t)$.

A.2 Filtration

Since we are interested in building dynamic models in which uncertainty is revealed over time, we need to build a proper **information structure**. Mathematically, this information structure will be constructed using a sequence of σ -algebras $\{\mathcal{F}_t\}_{t=0}^\infty$, called a **filtration**.

The σ -algebra \mathcal{F}_t contains all information available up to time t , i.e., information contained in the partial histories s^t . We will therefore also refer to \mathcal{F}_t as an **information set**. Remember that a σ -algebra is a collection of sets of paths. There is a **simple test** to decide whether a set of paths $F \subseteq \Omega$ belongs to the σ -algebra \mathcal{F}_t . We will have $F \in \mathcal{F}_t$ if, and only if, for every path $s^\infty \in \Omega$, we can decide, using information available up to time t only, whether $s^\infty \in F$ or $s^\infty \notin F$.

There is a straightforward procedure how to build \mathcal{F}_t . First, include in it all sets $F(s^t)$, i.e., sets whose partial history up to time t is s^t , for every s^t . Then apply rules from Defi-

inition A.1 to add all the other required sets. Observe that the set $F(s^t)$ certainly belongs to \mathcal{F}_t because we can decide, using information up to time t , whether partial history s^t has been realized or not. On the other hand, a set $F(s^{t+1})$ does not generically belong to \mathcal{F}_t because at time t we do not have information about the realized state s_{t+1} .

We now defined all the components of the space on which we construct stochastic processes:

- sample space Ω , containing all infinite-horizon paths s^∞ ,
- σ -algebra \mathcal{F} , containing all measurable sets on Ω ,
- filtration $\{\mathcal{F}_t\}_{t=0}^\infty$, defining the process of revealing information over time, and
- P , the probability measure over sets in \mathcal{F} .

Definition A.3. A *filtered probability space* is the quadruplet $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t=0}^\infty, P)$.

A.3 Random variables and stochastic processes

The state space \mathcal{S} and hence the sample space Ω can be abstract and inconvenient for mathematical operations. For instance, paths s^∞ (which are in our case elements of Ω) can be sequences of coin flips, weather patterns, periods of recessions and expansions, etc. This leads us to define the concept of random variables. A **random variable** is a mapping from the sample space to \mathbb{R}^n that provides some information about the path s^∞ . This will allow us to transfer probabilistic relationships from the abstract space Ω to a metric space with convenient properties, in which rules of calculus apply.

A casual definition of a random variable is simple: a random variable x is a function $x : \Omega \rightarrow \mathbb{R}^n$, i.e., a function that assigns to every path in Ω a real vector from \mathbb{R}^n . However, there are some important details. In particular, we want to make sure that the random variable preserves information structure, i.e., it does not reveal more information than what the underlying σ -algebra \mathcal{F} reveals.

In order to formalize this, we need to define another σ -algebra, this time on \mathbb{R}^n . A natural one will be the so-called **Borel σ -algebra** \mathcal{B} . This σ -algebra is constructed by first including into \mathcal{B} all open sets in \mathbb{R}^n and then again, as usual, repeating the application of the rules in Definition A.1. The sets in the Borel σ -algebra are called **Borel sets**. This σ -algebra is incredibly rich and contains almost all thinkable sets (go to Wikipedia and search for ‘non-measurable set’ to learn more).

We can now proceed with the formal definition of the random variable.

Definition A.4. A *random variable* is a function $x : \Omega \rightarrow \mathbb{R}^n$ that is measurable with respect to a given σ -algebra, say \mathcal{F} , i.e., for every Borel set $B \subseteq \mathbb{R}^n$, the preimage $x^{-1}(B)$ is in \mathcal{F} :

$$x^{-1}(B) = \{\omega \in \Omega : x(\omega) \in B\} \in \mathcal{F}.$$

We will denote x_t a random variable that is measurable with respect to \mathcal{F}_t .

For example, when paths s^∞ are infinite paths of coin flip realizations, then we could define $x_t(s^\infty) = 1$ if the state s_t from the path s^∞ is equal to 'heads', and 0 otherwise. Or, $x_t(s^\infty)$ could be the sum of all 'heads' in the partial history s^t of the path s^∞ . When there is no confusion given the context, we may also denote x the realization of the random variable.

It is instructive to carefully investigate what Definition A.4 states. Consider, for instance, a random variable x_t measurable with respect to the σ -algebra \mathcal{F}_t . Now take a set $F(s^t)$ of all paths s^∞ with common history s^t . Notice that this set cannot be further subdivided into two or more nontrivial smaller sets that also belong to \mathcal{F}_t . We say that $F(s^t)$ is an element of the **finest cover** of Ω constructed using sets from \mathcal{F}_t .

We want to argue that the random variable $x_t : \Omega \rightarrow \mathbb{R}^n$ that is measurable with respect to \mathcal{F}_t must be constant on $F(s^t)$, and that this property illustrates the notion of preservation of information.

To show this, take two paths $s^\infty, \tilde{s}^\infty \in F(s^t)$, i.e., two paths that share the common partial history s^t and assume that $x_t(s^\infty) \neq x_t(\tilde{s}^\infty)$. Since individual points in \mathbb{R}^n are Borel sets, then $x_t(s^\infty)$ and $x_t(\tilde{s}^\infty)$ are two distinct Borel sets. Hence, according to Definition A.4, the preimages of these sets (formally denoted $x_t^{-1}(x_t(s^\infty))$ and $x_t^{-1}(x_t(\tilde{s}^\infty))$), have to form two disjoint sets, both belonging to \mathcal{F}_t and both being nontrivial subsets of $F(s^t)$. But this is a contradiction, because $F(s^t)$ cannot be further subdivided into nontrivial subsets from \mathcal{F}_t .

Notice that this is an expected outcome. Imagine that the random variable x_t would indeed assign two different values $x_t(s^\infty) \neq x_t(\tilde{s}^\infty)$ to the paths $s^\infty, \tilde{s}^\infty \in F(s^t)$. This means that this random variable is able to distinguish s^∞ and \tilde{s}^∞ . However, these two paths have an identical partial history s^t and hence cannot be distinguished using information available at time t . Hence x_t would not be measurable with respect to \mathcal{F}_t .

To summarize, a random variable x_t measurable with respect to \mathcal{F}_t reveals some information available up to time t , but it cannot have 'precognition', i.e., it cannot reveal information that will only be learned after time t .

Since $x_t(s^\infty)$ is constant on the set of all paths s^∞ with a common partial history s^t , we will also use the notation $x_t(s^t)$ to denote the value of the random variable for an arbitrary $s^\infty \in F(s^t)$.

We can now move to the definition of a **stochastic process**.

Definition A.5. A *stochastic process* $\{x_t\}_{t=0}^\infty$ on the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t=0}^\infty, P)$ *adapted to* $\{\mathcal{F}_t\}_{t=0}^\infty$ is a sequence of random variables x_t , each of which is \mathcal{F}_t -measurable.

The sequence of random variables that forms the stochastic process $\{x_t\}_{t=0}^\infty$ therefore over time reveals progressively more information about the path s^∞ . Notice the following:

- $x_t(s^\infty)$ is a vector in \mathbb{R}^n that is the realized value of the stochastic process at time t on path s^∞ ,
- $x_t(\cdot)$, viewed as a function of $s^\infty \in \Omega$ for a fixed t , is a random variable measurable with respect to \mathcal{F}_t that assigns a vector in \mathbb{R}^n to each s^∞ ,

- $x. (s^\infty)$, viewed as a function of time ($x. (s^\infty) : \mathcal{T} \rightarrow \mathbb{R}^n$) for a fixed s^∞ is a deterministic function that assigns to every time $t \in \mathcal{T}$ a vector in \mathbb{R}^n , and is called the **path** of the stochastic process.

A.4 The Radon–Nikodým derivative

We are interested in ways how to relate two probability measures defined on the same sample space. This is done using a random variable called the **Radon–Nikodým derivative**. This random variable is utilized in many situations, for example to define conditional expectations, changes of measure or probability densities from distributions. We first need to define a property that relates two measures called absolute continuity.

Definition A.6. A probability measure \tilde{P} is **absolutely continuous** with respect to P if for every measurable set $F \in \mathcal{F}$, $P(F) = 0$ implies $\tilde{P}(F) = 0$. Measures \tilde{P} and P are **equivalent** if \tilde{P} is absolutely continuous with respect to P and P is absolutely continuous with respect to \tilde{P} .

In words, when \tilde{P} is absolutely continuous with respect to P , then \tilde{P} cannot assign a positive probability to a set to which P assigns a zero probability. For example, when \tilde{P} is the subjective probability measure of the agent, and P is the probability measure of the data-generating process, then the agent cannot assign a positive probability to a set of paths that has a zero probability according to the data-generating process.

We now state the formal theorem and then provide an intuitive construction.

Theorem 1.1 (Radon–Nikodým). Let \tilde{P} be absolutely continuous with respect to P . Then there exists a measurable function $h : \Omega \rightarrow \mathbb{R}_+$ such that

$$\tilde{P}(F) = \int_F h(\omega) dP(\omega). \quad (\text{A.1})$$

The function f is called the Radon–Nikodým derivative and commonly denoted as $h = d\tilde{P}/dP$.

The integral in (A.1) is taken in the sense of Remark A.1. Notice that Theorem 1.1 can be applied in two ways. Either we know P and \tilde{P} and the theorem gives us a way how to infer the function h , or we know P and the function h , and construct \tilde{P} using equation (A.1).

Also notice that h is a random variable measurable with respect to \mathcal{F} , assigning non-negative numbers to each path s^∞ . It will be useful to define **restrictions** of h to individual σ -algebras \mathcal{F}_t , denoted h_t . For a path $s^\infty \in F(s^t)$, define

$$h_t(s^\infty) \doteq \frac{\tilde{P}(s^t)}{P(s^t)} = \frac{\int_{\omega|s^t} h(\omega) d\tilde{P}(\omega)}{\int_{\omega|s^t} dP(\omega)}.$$

where $\omega|s^t$ denotes all paths $\omega \in \Omega$ that are continuations of the partial history s^t . Notice that $h_t(s^\infty)$ has the same value for all s^∞ with the same partial history s^t . Therefore, if there is no ambiguity, we will use the notation $h_t(s^\infty) = h(s^t)$, $\forall s^\infty \in F(s^t)$.

Since the Radon–Nikodým derivative can be interpreted as the ratio of two probability measures, h and its restrictions have to satisfy certain consistency requirements. In particular, we have the following relationship between restrictions of h to \mathcal{F}_t and \mathcal{F}_{t+j} :

$$\begin{aligned} h(s^t) &= \frac{\tilde{P}(s^t)}{P(s^t)} = \frac{\sum_{s^{t+j}|s^t} \tilde{P}(s^{t+j})}{P(s^t)} = \sum_{s^{t+j}|s^t} h(s^{t+j}) \frac{P(s^{t+j})}{P(s^t)} \\ &= \sum_{s^{t+j}|s^t} h(s^{t+j}) P(s^{t+j}|s^t). \end{aligned}$$

In line with these definitions, we can also define the restriction of the Radon–Nikodým derivative to conditional probabilities. In particular, let s^{t+j} be a continuation of s^t and $s^\infty \in F(s^{t+j})$. Then notice that

$$\tilde{P}(s^{t+j}|s^t) = \frac{\tilde{P}(s^{t+j})}{\tilde{P}(s^t)} = \frac{h(s^{t+j}) P(s^{t+j})}{h(s^t) P(s^t)} = \frac{h(s^{t+j})}{h(s^t)} P(s^{t+j}|s^t)$$

and therefore we can define

$$h_{t+j}(s^\infty|s^t) \doteq h(s^{t+j}|s^t) \doteq \frac{\tilde{P}(s^{t+j}|s^t)}{P(s^{t+j}|s^t)} = \frac{h(s^{t+j})}{h(s^t)}.$$

A.5 Expectations

Before we define the expectations operator for a random variable x , we define a probability measure μ_x that will represent the distribution of x . This distribution will allow us to compute moments of x by integrating over \mathbb{R}^n , rather than over the abstract set of paths Ω . We will do so by taking every Borel set $B \subset \mathbb{R}^n$, and obtaining its preimage $x^{-1}(B) \subset \Omega$. Since this preimage is measurable with respect to the relevant σ -algebra, i.e., $x^{-1}(B) \in \mathcal{F}$, it has been assigned a probability $P(x^{-1}(B))$. We then assign the same probability to the Borel set B , and denote it $\mu_x(B)$. Formally:

Definition A.7. Associated with the random variable $x : \Omega \rightarrow \mathbb{R}^n$ is the probability measure that, to each Borel set $B \subset \mathbb{R}^n$, assigns the probability

$$\mu_x(B) = P(x^{-1}(B)). \quad (\text{A.2})$$

μ_x is called the **distribution** of x . If $x : \Omega \rightarrow \mathbb{R}$, then the **cumulative distribution function** of x is the function $G_x : \mathbb{R} \rightarrow [0, 1]$

$$G_x(x) = \mu_x((-\infty, x])$$

Further, the function $g_x : \mathbb{R} \rightarrow \mathbb{R}_+$, if it exists, defined as

$$g_x(x) = \frac{\partial}{\partial x} G_x(x)$$

is called the *probability density function*.

The definition of the cumulative distribution function and probability density function can be extended to the multivariate case as well. We can now define the expectation of x , either by integrating over Ω using P , or integrating over \mathbb{R}^n using μ_x .

Definition A.8. *The expectation of a random variable is*

$$E[x] = \int_{\Omega} x(\omega) dP(\omega) = \int_{\mathbb{R}^n} x d\mu_x(x).$$

Observe that the two integrals are over different spaces, and use different probability measures but by equation (A.2), they must yield the same results. Finally, we can provide a simple definition of the **conditional expectation** conditional on an information set (σ -algebra) \mathcal{F}_t as follows (there is a more formal and more general definition that uses the Radon–Nikodým derivative).

Definition A.9. *The conditional expectations operator for a random variable x measurable with respect to \mathcal{F} , conditional on information set \mathcal{F}_t , is a function*

$$E[x | \mathcal{F}_t] : \Omega \rightarrow \mathbb{R}^n$$

such that for any $s^\infty \in \Omega$ with partial history s^t , we have

$$E[x | \mathcal{F}_t](s^\infty) = \frac{\int_{F(s^t)} x(\omega) dP(\omega)}{P(s^t)} = \int_F x(\omega) dP(\omega | s^t).$$

This expression simplifies if we know that $x = x_{t+j}$, i.e. it is measurable with respect to the information set \mathcal{F}_{t+j} . Then

$$E[x_{t+j} | \mathcal{F}_t](s^\infty) = \sum_{s^{t+j} | s^t} x_{t+j}(s^{t+j}) P(s^{t+j} | s^t). \quad (\text{A.3})$$

Notice that any x_u measurable with respect to \mathcal{F}_u for any $u \leq t$ is constant on any $F(s^t)$, so that $E[x_u | \mathcal{F}_t](s^\infty) = x_u(s^\infty)$.

Remark A.2. *When there is no ambiguity, we will use the notation $E_t[x]$ for the conditional expectation, and P_t for the conditional probability measure conditional on an event from \mathcal{F}_t . Similarly, for an $s^\infty \in F(s^t)$, we will write $E[x_{t+j} | s^t]$ instead of $E[x_{t+j} | \mathcal{F}_t](s^\infty)$.*

From (A.3), we can also define conditional expectations conditional on a different set of paths than $F(s^t)$. In particular, for a conditioning set $H \in \mathcal{F}_t$, we can define the conditional

expectation as

$$\begin{aligned}
 E[x_{t+j}|H] &\doteq \sum_{F(s^t) \in H} E[x_{t+j}|s^t] P(s^t|H) = \frac{1}{P(H)} \sum_{F(s^t) \in H} E[x_{t+j}|s^t] P(s^t) \\
 &= \frac{1}{P(H)} \sum_{F(s^{t+j}) \in H} x_{t+j}(s^{t+j}) P(s^{t+j}) \\
 &= \sum_{F(s^{t+j}) \in H} x_{t+j}(s^{t+j}) P(s^{t+j}|H)
 \end{aligned}$$

where the third equality follows from substituting in the expression from (A.3). This result will be particularly useful when we condition on realizations of a random variable, rather than partial histories s^t .

A.6 Martingales

An important class of processes are processes that on average neither grow nor decay, called martingales.

Definition A.10. An n -dimensional process $\{x_t\}_{t=0}^{\infty}$ on $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t=0}^{\infty}, P)$ is a *martingale* with respect to the filtration $\{\mathcal{F}_t\}$ and the probability measure P if:

1. x_t is \mathcal{F}_t -measurable (i.e., the process $\{x_t\}_{t=0}^{\infty}$ is adapted),
2. $E[|x_t|] < \infty$ for all $t \in \mathcal{T}$,
3. $E[x_s|\mathcal{F}_t] = x_t$ for all $s \geq t$.

A *supermartingale* is a process for which condition 3. is replaced by $E[x_s|\mathcal{F}_t] \leq x_t$, and a *submartingale* is a process for which condition 3. is replaced by $E[x_s|\mathcal{F}_t] \geq x_t$.

The martingale property is always defined with respect to a particular filtration and a probability measure. A process that is a martingale under a given filtration and probability measure will typically not be a martingale under a different filtration or a different probability measure.

A.7 The Markov property

“We may regard the present state of the universe as the effect of its past and the cause of its future.” Marquis de Laplace

We want to impose several properties on stochastic processes that will interest us: **Markovianity**, **time-invariance**, **stationarity** and **ergodicity**. The following properties are important when we want to compare predictions of our models with time series observations in the data.

Definition A.11. A process $\{x_t\}_{t=0}^{\infty}$ on $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t=0}^{\infty}, P)$ is said to be **Markov** if $\forall t, s \geq 0$ and for every Borel set $B \in \mathcal{B}$

$$P(x_{t+s} \in B | \mathcal{F}_t) = P(x_{t+s} \in B | x_t). \quad (\text{A.4})$$

A Markov process is **time-invariant** if the conditional distribution $P(x_{t+s} \in B | x_t = x)$ only depends on x and not on t .

A process (not necessarily Markov) is **stationary** if its joint distribution does not depend on time, i.e.,

$$\forall k, h \in \mathbb{N}, \forall (t_1, \dots, t_k) \in \mathbb{N}^k, \forall B_{t_i} \in \mathcal{B}$$

we have

$$P(x_{t_1} \in B_{t_1}, \dots, x_{t_k} \in B_{t_k}) = P(x_{t_1+h} \in B_{t_1}, \dots, x_{t_k+h} \in B_{t_k}). \quad (\text{A.5})$$

Definition A.12. A **stationary distribution** of a Markov process is the marginal distribution over the state x_t such that $\forall t, h \in \mathbb{N}$

$$P(x_t \in B) = P(x_{t+h} \in B). \quad (\text{A.6})$$

We postpone an exact definition of ergodicity until we discuss Markov chains. The above definitions can be verbally interpreted as follows.

- A process $\{x_t\}_{t=0}^{\infty}$ is **Markov** if the distribution of its future realizations conditional on all information available up to time t is the same as if we conditioned on x_t only. Hence x_t summarizes all information needed to describe the future distribution of the process, and therefore the realization of x_t can serve as a **state vector**.

When the information set is generated only by the history of realizations of $\{x_t\}_{t=0}^{\infty}$, then we can write (A.4) as

$$P(x_{t+s} \in B | x_t, x_{t-1}, x_{t-2}, \dots) = P(x_{t+s} \in B | x_t).$$

- A Markov process is **time-invariant** when its transition density (i.e., the conditional distribution) does not depend on calendar time.
- A (general) process is **stationary** when the joint distribution of its realizations (not conditional!) stays the same over time. Notice that for a time-invariant **Markov** process to be stationary, it suffices to show property (A.6).
- A process is **ergodic** when its stochastic properties can be deduced from one sufficiently long (in principle infinite) observed path of realizations.

Ljungqvist and Sargent (2018), Chapter 2, discuss extensive applications of the above properties in the context of Markov chains.

The Markov property is central to our modeling approach based on dynamic programming. Since the Markov state summarizes all information available at time t relevant for the future evolution of the stochastic process, it will play the role of a state variable in the dynamic recursive program. Finding the right vector of state variables that leads to a

Markov representation of an economic problem is often nontrivial (“*finding the state is an art*”) but macroeconomists found many creative ways how to set up recursive representations using suitable state variables for problems that have been previously considered intractable.

A.8 Macroeconomic models and parameters

At this stage, we can return to the general discussion from the introduction to this chapter. We built the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t=0}^\infty, P)$ and defined a Markov process $\{x_t\}_{t=0}^\infty$, and we will now associate these Markov processes with solutions to macroeconomic models. A **model** is therefore a probability distribution over $\{x_t\}_{t=0}^\infty$.

Macroeconomic models that we are interested in are parameterized by a vector of **structural parameters** $\theta \in \Theta$ that are presumed to be known to the agents in the model and the theorist building the model takes them as given, but are unknown to the econometrician. These parameters include preference and technology parameters, parameters describing the information structure, market interaction, or the distribution of structural shocks. Each θ thus represents an alternative macroeconomic model, and hence an alternative probability distribution $P(\theta)$ that determines the distribution of the stochastic process $\{x_t\}_{t=0}^\infty$. The role of the econometrician is to collect data that are informative about $\{x_t\}_{t=0}^\infty$ and infer which θ from the parameter set Θ is the most likely candidate for the ‘true’ model that generated these data.

More formally, we can imagine that the model generates predictions for data y_t that are functions of the current state x_t given by the observation or **measurement equation**

$$y_t = f(x_t, v_t)$$

where $\{v_t\}_{t=0}^\infty$ represents noise or measurement errors. We thus have a probability distribution $P(y^t|\theta)$ for the data $y^t = (y_0, \dots, y_t)$. The **direct problem** (faced by a macroeconomist) consists of solving the macroeconomic model for a given parameter vector θ and obtaining the distribution $P(y^t|\theta)$. The **inverse problem** (faced by an econometrician) is to take observed data $\hat{y}^t = (\hat{y}_0, \dots, \hat{y}_t)$ and infer which probability distribution $P(y^t|\theta)$ from the family indexed by $\theta \in \Theta$ do the data come from.

We will briefly discuss some methods how to approach this inverse problem in this chapter. In particular, we will discuss the **maximum likelihood** approach, **generalized method of moments** due to Hansen (1982), a **Bayesian** approach to estimation, as well as **calibration** of macroeconomic models. It is important to stress that there is no uniquely preferred approach here, and there are substantial differences of opinion concerning the appropriate balance between emphasis put on theory and modeling on one side, and empirics and data. In reality, the right approach is driven by the economic problem itself.

An excellent introduction into these concepts and various perspectives that relate to macroeconomic modeling is a series of papers by Hansen and Heckman (1996), Sims (1996), and Kydland and Prescott (1996) published as a symposium in Journal of Economic Perspectives.

Bibliography

- Albrecht, James (2011) "The 2010 Nobel Memorial Prize in Search Theory," Mimeo, Georgetown University. 75
- Arrow, Kenneth J. (1964) "The Role of Securities in the Optimal Allocation of Risk-Bearing," *Review of Economic Studies*, 31 (2), 91–96, Translation from the French original, published in 1953. 27, 102
- (1965) *Aspects of the Theory of Risk Bearing*: Yrjö Jahnssonin Säätiö, Helsinki. 114
- Bellman, Richard E. (1957) *Dynamic Programming*: Princeton University Press, Princeton, NJ. 59
- Bewley, Truman F. (1980) "The Optimum Quantity of Money," in Kareken, John H. and Neil Wallace eds. *Models in Monetary Economics*: Federal Reserve Bank of Minneapolis, Minneapolis, MN. 36
- Bhandari, Anmol, Jaroslav Borovička, and Paul Ho (2019) "Survey Data and Subjective Beliefs in Business Cycle Models," Federal Reserve Bank of Richmond Working Paper No. 19-14. 93, 107, 115
- Billingsley, Patrick (1979) *Probability and Measure*: John Wiley & Sons, Inc. 12
- Bishop, Christopher M. (2006) *Pattern Recognition and Machine Learning*: Springer. 1
- Black, Fischer and Myron Scholes (1973) "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81 (3), 637–654. 133, 152, 155, 156, 158, 161, 172
- Blanchard, Olivier Jean and Charles M. Kahn (1980) "The Solution of Linear Difference Models under Rational Expectations," *Econometrica*, 48 (5), 1305–1312. 93, 118
- Borovička, Jaroslav and Lars Peter Hansen (2014) "Examining Macroeconomic Models through the Lens of Asset Pricing," *Journal of Econometrics*, 183 (1), 67–90. 93, 107, 115
- Borovička, Jaroslav and John Stachurski (2021) "Stability of Equilibrium Asset Pricing Models: A Necessary and Sufficient Condition," *Journal of Economic Theory*, 193, 105227. 47, 78
- Breeden, Douglas T. and Robert H. Litzenberger (1978) "Prices of State-Contingent Claims Implicit in Option Prices," *Journal of Business*, 51 (4), 621–651. 159

- Burdett, Kenneth and Kenneth L. Judd (1983) "Equilibrium Price Dispersion," *Econometrica*, 51 (4), 955–969. 74
- Burdett, Kenneth and Dale T. Mortensen (1998) "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 39 (2), 257–273. 75
- Cagan, Phillip D. (1956) "The Monetary Dynamics of Hyperinflation," in Friedman, Milton ed. *Studies in the Quantity Theory of Money*, 25–117: University of Chicago Press. 124, 125, 126
- Campbell, John Y. and Robert J. Shiller (1988) "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1, 195–228. 93, 106, 109
- Candler, Graham V. (2001) "Finite-Difference Methods for Continuous-Time Dynamic Programming," in *Computational Methods for the Study of Dynamic Economies*, Chap. 8, 172–194: Oxford University Press. 133, 171, 172
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (2005) "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113 (1), 1–45. 119, 128, 131
- Cochrane, John H. (2005) *Asset Pricing*: Princeton University Press, Revised Edition. 1, 3
- Cox, John C. and Stephen A. Ross (1976) "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics*, 3 (1–2), 145–166. 158
- Cox, John C., Stephen A. Ross, and Mark Rubinstein (1979) "Option Pricing: A Simplified Approach," *Journal of Financial Economics*, 7 (3), 229–263. 133, 159, 172, 174
- Debreu, Gerard (1959) *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*: Yale University Press, New Haven and London. 27, 102
- Diamond, Peter A. (1971) "A Model of Price Adjustment," *Journal of Economic Theory*, 3 (2), 156–168. 74
- Diamond, Peter A. and Joseph E. Stiglitz (1974) "Increases in Risk and in Risk Aversion," *Journal of Economic Theory*, 8 (3), 337–360. 64
- Duffie, Darrell (2001) *Dynamic Asset Pricing Theory*: Princeton University Press, 3rd edition. 1, 53, 133
- Evans, George W. and Seppo Honkapohja (2001) *Learning and Expectations in Macroeconomics*: Princeton University Press, Princeton, NJ. 177, 191
- Evans, George W., Seppo Honkapohja, and Thomas J. Sargent (2005) "An Interview with Thomas J. Sargent," *Macroeconomic Dynamics*, 9 (4), 561–583. 191
- Fisher, Irving (1896) "Appreciation and Interest," *Publications of the American Economic Association*, 11 (4), 1–98. 102

-
- (1906) *The Nature of Capital and Income*: Macmillan, New York. 102
- (1930) *The Theory of Interest*: Macmillan, New York. 102
- Friedman, Milton (1957) *A Theory of the Consumption Function*: Princeton University Press, Princeton. 124, 125, 126
- Frisch, Ragnar (1933) "Propagation Problems and Impulse Problems in Dynamic Economics," in *Economic Essays in Honour of Gustav Cassel*, 171–205: Allen and Unwin. 96, 99
- Galí, Jordi (2008) *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*: Princeton University Press. 1
- Gil, Amparo, Javier Segura, and Nico M. Temme (2007) *Numerical Methods for Special Functions*: SIAM, Society for Industrial and Applied Mathematics, Philadelphia. 80
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville (2016) *Deep Learning*: MIT Press, <http://www.deeplearningbook.org>. 1
- Hansen, Lars Peter (1982) "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50 (4), 1029–1054. 3, 16, 37, 41, 131, 203
- (2001) "Method of Moments," in Smelser, Neil J. and Paul B. Baltes eds. *International Encyclopedia of the Social & Behavioral Sciences*, 9743–9751: Elsevier. 41
- (2008) "Generalized Method of Moments Estimation," in Durlauf, Steven N. and Lawrence Blume eds. *The New Palgrave Dictionary of Economics*: Palgrave Macmillan, London. 3, 41
- (2014) "Nobel Lecture: Uncertainty Outside and Inside Economic Models," *Journal of Political Economy*, 122 (5), 945–987. 37
- Hansen, Lars Peter, John C. Heaton, and Amir Yaron (1996) "Finite-Sample Properties of Some Alternative GMM Estimators," *Journal of Business & Economic Statistics*, 14 (3), 262–280. 42, 43
- Hansen, Lars Peter and James J. Heckman (1996) "The Empirical Foundations of Calibration," *Journal of Economic Perspectives*, 10 (1), 87–104. 203
- Hansen, Lars Peter and Kenneth J. Singleton (1982) "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 50 (5), 1269–1286. 3, 37, 39
- (1983) "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," *Journal of Political Economy*, 91 (2), 249–265. 3, 44, 50, 131
- Harrison, J. Michael and David M. Kreps (1979) "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory*, 20 (3), 381–408. 52

- Harrison, J. Michael and Stanley R. Pliska (1981) "Martingales and Stochastic Integrals in the Theory of Continuous Trading," *Stochastic Processes and their Applications*, 11 (3), 215–260. 52
- Heston, Steven L. (1993) "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *Review of Financial Studies*, 6 (2), 327–343. 158
- Hicks, John Richard (1939) *Value and Capital: An Inquiry into Some Fundamental Principles of Economic Theory*: Clarendon Press, Oxford. 102
- Holmes, Mark H. (1995) *Introduction to Perturbation Methods*: Springer. 93, 107
- (2007) *Introduction to Numerical Methods in Differential Equations*: Springer. 1, 133
- Howard, Ronald A. (1960) *Dynamic Programming and Markov Processes*: MIT Press. 104
- Itô, Kiyosi (1951) "On a Formula Concerning Stochastic Differentials," *Nagoya Mathematical Journal*, 3, 55–65. 148
- Jovanovic, Boyan (1979) "Job Matching and the Theory of Turnover," *Journal of Political Economy*, 87 (5), 972–990. 121, 126
- Judd, Kenneth L. (1998) *Numerical Methods in Economics*: The MIT Press, Cambridge, MA. 1, 78, 93, 107, 133
- Kac, Marc (1949) "On Distributions of Certain Wiener Functionals," *Transactions of the American Mathematical Society*, 65 (1), 1–13. 151
- Kalman, Rudolph Emil (1960) "New Approach to Linear Filtering and Prediction Problems," *Transactions of the ASME—Journal of Basic Engineering*, 82 (Series D), 35–45. 121, 125, 137
- Kalman, Rudolph Emil and Richard Snowden Bucy (1961) "New Results in Linear Filtering and Prediction Theory," *Journal of Basic Engineering*, 83 (1), 95–108. 137
- Kamihigashi, Takashi (2003) "Necessity of Transversality Conditions for Stochastic Problems," *Journal of Economic Theory*, 109 (1), 140–149. 28
- (2005) "Necessity of the Transversality Condition for Stochastic Models with Bounded or CRRA Utility," *Journal of Economic Dynamics and Control*, 29 (8), 1313–1329. 28
- Kreps, David M. (1981) "Arbitrage and Equilibrium in Economies with Infinitely Many Commodities," *Journal of Mathematical Economics*, 8 (1), 15–35. 52
- Kydland, Finn E. and Edward C. Prescott (1982) "Time to Build and Aggregate Fluctuations," *Econometrica*, 50 (6), 1345–1370. 96
- (1996) "The Computational Experiment: An Econometric Tool," *Journal of Economic Perspectives*, 10 (1), 69–85. 203

- Li, Xiang, Wenhao Yang, Jiadong Liang, Zhihua Zhang, and Michael I. Jordan (2023) "A Statistical Analysis of Polyak–Ruppert Averaged Q-learning," February, <https://arxiv.org/abs/2112.14582>. 182
- Ljungqvist, Lars and Thomas J. Sargent (2012) *Recursive Macroeconomic Theory*: The MIT Press, 3rd edition. 131
- (2018) *Recursive Macroeconomic Theory*: The MIT Press, 4th edition. 1, 3, 33, 55, 74, 93, 104, 118, 121, 202
- Lombardo, Giovanni (2010) "On Approximating DSGE Models by Series Expansions," ECB Working paper No. 1264. 93, 107
- Lucas, Jr., Robert E. (1972) "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 4 (2), 103–124. 34
- (1978) "Asset Prices in an Exchange Economy," *Econometrica*, 46 (6), 1429–1445. 3, 4, 32, 43
- McCall, John (1970) "Economics of Information and Job Search," *Quarterly Journal of Economics*, 84 (1), 113–126. ii, 24, 55, 60, 61, 62, 63, 65, 66, 72, 75, 177, 179, 183, 184, 187
- Mehra, Rajnish and Edward C. Prescott (1985) "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15 (2), 145–161. 3, 4, 43, 48, 49, 50, 51
- Merton, Robert C. (1973) "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, 4 (1), 141–183. 133, 152
- (1977) "Option Pricing When Underlying Stock Returns Are Discontinuous," *Journal of Financial Economics*, 3 (1–2), 125–144. 158
- Muth, John F. (1960) "Optimal Properties of Exponentially Weighted Forecasts," *Journal of the American Statistical Association*, 55 (290), 299–306. 121, 125
- (1961) "Rational Expectations and the Theory of Price Movements," *Econometrica*, 29 (3), 315–335. 34, 125
- Nelson, Daniel B. and Krishna Ramaswamy (1990) "Simple Binomial Processes as Diffusion Approximations in Financial Models," *Review of Financial Studies*, 3 (3), 393–430. 172
- von Neumann, John and Oskar Morgenstern (1947) *Theory of Games and Economic Behavior*: Princeton University Press, 2nd edition. 26, 29
- Newey, Whitney K. and Kenneth D. West (1987) "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55 (3), 703–708. 42
- Øksendal, Bernt (2007) *Stochastic Differential Equations*: Springer Verlag, 6th edition. 133, 140, 145, 151

- Pratt, John W. (1965) "Risk Aversion in the Small and in the Large," *Econometrica*, 32 (1–2), 122–136. 114
- Ross, Stephen A. (1976) "The Arbitrage Theory of Capital Asset Pricing," *Journal of Economic Theory*, 13 (3), 341–360. 133
- (1978) "A Simple Approach to the Valuation of Risky Streams," *Journal of Business*, 51 (3), 453–475. 158
- Rothschild, Michael and Joseph E. Stiglitz (1970) "Increasing Risk: I. A Definition," *Journal of Economic Theory*, 2 (3), 225–243. 63, 64
- Samuelson, Paul A. (1939) "Interactions between the Multiplier Analysis and the Principle of Acceleration," *Review of Economics and Statistics*, 21 (2), 75–78. 96
- Santos, Manuel S. and Michael Woodford (1997) "Rational Asset Pricing Bubbles," *Econometrica*, 65 (1), 19–57. 36
- Sargent, Thomas J. (1993) *Bounded Rationality in Macroeconomics*: Clarendon Press, Oxford. 177, 191
- (1999) *The Conquest of American Inflation*: Princeton University Press, Princeton, NJ. 177, 191
- Savage, Leonard J. (1954) *The Foundations of Statistics*: Wiley, New York. 29, 33
- Sims, Christopher A. (1996) "Macroeconomics and Methodology," *Journal of Economic Perspectives*, 10 (1), 105–120. 203
- (2002) "Solving Rational Expectations Models," *Computational Economics*, 20 (1–2), 1–20. 93, 118
- Slutsky, Eugen (1927) "The Summation of Random Causes as the Source of Cyclic Processes," in *Problems of Economic Conditions*, 3, Chap. 1: The Conjecture Institute, Moscow. 95, 96
- (1937) "The Summation of Random Causes as the Source of Cyclic Processes," *Econometrica*, 5 (2), 105–146. 95
- Stokey, Nancy L., Robert E. Lucas, and Edward C. Prescott (1989) *Recursive Methods in Economic Dynamics*: Harvard University Press. 1, 28, 30, 31
- Sutton, Richard S. and Andrew G. Barto (2018) *Reinforcement Learning: An Introduction*: MIT Press, Cambridge, MA, 2nd edition. 1, 177, 182, 187, 191
- Thomas, James W. (1995) *Numerical Partial Differential Equations: Finite Difference Methods*: Springer Verlag New York. 133
- Tirole, Jean (1985) "Asset Bubbles and Overlapping Generations," *Econometrica*, 53 (5), 1071–1100. 36

Yule, George Udny (1927) "On a Method of Investigating Periodicities in Disturbed Series, with Special Reference to Wolfer's Sunspot Numbers," *Philosophical Transactions of the Royal Society of London. Series A*, 226 (636–646), 267–298. 96

Index

- adaptive expectations, 124
- Arrow–Debreu prices, 102

- backward induction, 71, 76
- Bayes theorem, 131, 195
- Bayesian estimation, 130, 203
- beliefs, 57
- Bellman equation, 58
- Bellman operator, 59
- Black–Scholes model, 152
- Brownian motion, 138
- bubble, 23

- calibration, 203
- cash flow, 102
- certainty equivalence principle, 105
- cumulative distribution function, 199

- data-generating process, 57
- decision rule, 58
- Diamond paradox, 74
- difference equation
 - stochastic linear, 94
- discounting
 - geometric, 56
- discrete Lyapunov equation, 98, 102
- doubling algorithm, 98
- dynamic programming, 59

- eigenvalue
 - Perron–Frobenius, 5
- eigenvector, 5
 - base, 7
- elementary process, 142
- equity premium puzzle, 50
- ergodic theorem, 10, 12
- ergodicity, 10
 - Markov chain, 11

- Euler equation, 28
- event, 194

- Feynman–Kac formula, 151
- filtering problem
 - Markov chain, 18
- filtration, 195
- fixed point, 59

- generalized method of moments, 16, 131, 203
- Golub–Welsch algorithm, 86
- Gram–Schmidt orthogonalization, 85
- Gram–Schmidt orthogonalization, 122

- \mathcal{H}^2 class, 145
- $\mathcal{H}^2(X)$ class, 147
- hidden Markov model, 17

- impulse response function, 100
 - linear, 100
 - matching, 128, 131
- inference problem, 15
- information set, 57, 195
- invariant function, 10
- Itô diffusion, 146
- Itô integral, 145
 - for simple processes, 142
- Itô process, 146
 - drift and volatility, 146
 - per partes formula, 149
- Itô’s lemma, 148
 - multivariate, 148

- Kalman filter, 121
 - innovations representation, 127
 - Kalman gain, 123
- Kalman smoother, 124

- \mathcal{L}^1 class, 145

- \mathcal{L}^2 class, 145
- $\mathcal{L}(X)$ class, 147
- Legendre polynomial, 89
- linear model
 - multivariate stochastic, 94
- Markov chain, 4
 - hidden, 17
 - irreducible, 7
- Markov process, 193, 202
 - Brownian motion, 139
- martingale, 11, 201
 - Brownian motion, 139
- martingale convergence theorem, 11
- martingale difference sequence, 95
- maximum likelihood estimator, 15, 203
- MCMC methods, 131
- mean-preserving spread, 63
- measurability
 - constraints, 27
- measurable set, 195
- measurement equation, 18, 95, 203
- method of moments, 16
- mixing condition, 5, 12
- model, 203
- moment conditions, 16
- moving-average representation, 100
- multiplier-accelerator model, 96
- no-Ponzi condition, 30
- observation equation, *see* measurement equation
- option pricing, 152
 - dynamic hedging, 155
 - implied volatility, 158
- option value, 62, 68
- outcome, 194
- parameter
 - structural, 203
- partial history, 193
 - continuation, 194
- partial identification, 38
- partition
 - interval, 139, 141
- path, 193
- Perron–Frobenius theorem, 5
- policy function, 59
- principle of optimality, 59, 68
- probability
 - conditional, 195
 - unconditional, 194
- probability density function, 200
- probability measure, 194
 - absolutely continuous, 198
 - equivalent, 198
- probability space
 - filtered, 196
- put-call parity, 157
- quadrature rule, 78
 - Gaussian quadrature, 79
- Radon–Nikodým derivative, 198
- random variable, 196
 - conditional expectation, 200
 - distribution, 199
 - expectation, 200
- rational expectations, 57
- reservation wage, 60
- Riccati equation, 124
- Riemann–Stieltjes integral, 141
- risk premium, 103
- σ -algebra, 194
- sample space, 193
- securities
 - derivative, 154
- sequence problem, 56
- σ -algebra
 - Borel, 196
- simple process, *see* elementary process
- single-crossing property, 63
- Slutsky–Yule effect, 96
- solution
 - bubble, 23
 - fundamental, 22
 - Markov, 23
- state, 193

- state variable, 57, 59
 - finding the state is an art, 57
- state-space representation, 95
- stationarity
 - covariance, 97
- stochastic discount factor, 29, 102
- stochastic process, 197
 - predictable, 135
- strategy, 57
- Stratonovich integral, 144
- submartingale, 201
- successive approximation, 76
- supermartingale, 201

- time-invariance, 58
- transversality condition, 30

- upwind differencing, 169

- value function, 56
- value function iteration, 76, 105
 - Howard improvement algorithm, 104
 - policy function iteration, 104
- variation process
 - cross-variation, 140
 - quadratic, 140
 - total, 139
- vector autoregression, 94
- verification theorem, 59

- welfare, 65
 - ex-ante, 67
- white noise, 95