

Lecture Notes on Stochastic Processes in Economics

Srikanth Pai, Asst. Professor, Madras School of Economics, Chennai

October 2, 2025

Contents

1	Discrete Random Variables and Sigma-Algebras	3
1.1	Motivation	3
1.2	Sigma-Algebras	3
1.3	Examples	3
1.4	Measurable Functions	7
2	Probability Measures and Expectation	8
2.1	Probability Measure	8
2.2	Expectation	9
2.3	Example	9
2.4	Discussion	9
2.5	Exercises	10
3	Markov Decision Processes	10
3.1	Motivation	10
3.2	Definitions	11
3.3	Examples: The McCall Job Search Problem	12
3.4	Discussion	13
3.5	Exercises	14
4	Conditional Expectation	15
4.1	Motivation	15
4.2	Conditional Expectation with respect to Events	16
4.3	Conditional expectation via partitions: an economic perspective	18
4.4	Conditional Expectation: Defining property	19
4.5	Conditional Expectation with respect to a Random Variable	21
4.6	Properties of conditional expectation	22
4.7	Applications	24
4.8	Exercises	24

5	Martingales	26
5.1	Filtered Probability Spaces and Adapted Processes	26
5.2	Martingales, Submartingales, and Supermartingales	26
5.3	Basic Examples	27
5.4	Previsible Processes and Martingale Transforms	28
5.5	Stopping Times and Stopped Processes	29
5.5.1	Applications of Optional Sampling	31
5.5.2	Important Non-Example: the doubling (“martingale”) strategy	32
5.6	Monkeys typing ABRACADABRA	33
5.7	Markov Chains and Martingales	34
5.7.1	Explicit Harmonic Examples and Optional stopping calculations	35
5.8	Exercises	37
6	Martingales in Financial Markets	39
6.1	Risk Preferences and Utility	39
6.1.1	Definitions	40
6.2	Pricing a Risky Asset	42
6.2.1	Examples and Special Cases	43
6.3	Equivalent Measures and Risk–Neutral Measures	47
6.3.1	Interpretation of equivalent measures	47
6.3.2	Equivalent and risk–neutral measures in asset pricing	49
6.4	Fundamental Theorem of Asset Pricing	51
6.4.1	Convex Geometry Preliminaries	51
6.4.2	Trading Setup	52
6.4.3	The Fundamental Theorem	54
6.4.4	The theorem	55
6.4.5	Applications	57
6.4.6	Numerical Exercises	58

Preface

These notes collect material on stochastic processes relevant to economics, particularly dynamic programming and recursive methods in macroeconomics. The style is rigorous but oriented toward applications, with emphasis on motivation, worked examples, and exercises.

The notes contain a lot of remarks pertaining to economic applications or philosophy. These remarks can be safely skipped without affecting the flow of the chapter. Theorem proofs are usually restricted to the special case of partition sigma algebras and discrete random variables. The theory for these cases is covered carefully in the first chapter.

1 Discrete Random Variables and Sigma-Algebras

1.1 Motivation

A central idea in probability theory is measurability. At first sight this seems like an unnecessary technicality, but it becomes indispensable once we view probability as a language for describing information. In economics and finance, uncertainty is not only about which outcome occurs but also about what agents know. The structure that encodes what is knowable is a sigma-algebra, a collection of events that can be distinguished. Measure theory therefore provides the rigorous connection between random variables, which assign numerical values to states of the world, and the partitions of information that agents actually perceive.

One way to build intuition is to imagine an economy subject to shocks. Nature selects a state from some set of possibilities, but an agent does not observe this state directly. Instead the agent learns only which cell of a partition the state belongs to. For instance, the agent may only be able to tell whether productivity is high or low, without knowing the precise underlying disturbance. The partition of states generates a sigma-algebra, and random variables that are measurable with respect to this sigma-algebra are precisely those variables whose values depend only on the coarsened information conveyed by the partition. In this way, conditional expectation can be understood as averaging given what is known, and the formalism of sigma-algebras becomes the natural language for linking probability with economic decision-making.

1.2 Sigma-Algebras

A sigma-algebra on a set Ω is a collection \mathcal{F} of subsets of Ω that is closed under complements and countable unions, and that contains the whole set Ω itself. The elements of \mathcal{F} are called measurable sets or events. Intuitively, \mathcal{F} represents the events about which we are able to speak in probabilistic terms, so that probabilities can be consistently assigned.

Given any family of subsets of Ω , one can ask for the smallest sigma-algebra that contains them. This is called the sigma-algebra generated by the family.

1.3 Examples

The construction of sigma-algebras is easiest to understand through simple cases.

Suppose first that we start with a single subset $A \subseteq \Omega$. The sigma-algebra generated by A must contain Ω and \emptyset , and it must also contain A and its complement A^c . In fact no further sets are needed, so we obtain

$$\sigma(\{A\}) = \{\emptyset, A, A^c, \Omega\}.$$

If instead we start with two disjoint subsets $A, B \subseteq \Omega$, the generated sigma-algebra must contain A and B as well as their complements. Because sigma-algebras are closed under unions, the set

$A \cup B$ must also appear. Thus the sigma-algebra consists of the eight sets

$$\emptyset, A, B, A \cup B, A^c, B^c, (A \cup B)^c, \Omega.$$

This collection is precisely all unions of the partition $\{A, B, (A \cup B)^c\}$.

More generally, if Ω is partitioned into finitely many disjoint cells A_1, \dots, A_n , then the sigma-algebra generated by this partition is exactly the collection of all possible unions of the cells. As proved below, this sigma-algebra is the smallest one containing the partition. In this way the sigma-algebra encodes the information that distinguishes only which cell of the partition contains the true state of the world.

Theorem 1.1. *Let Ω be a set and let $\mathcal{P} = \{A_1, \dots, A_n\}$ be a finite partition of Ω ; that is, each A_i is nonempty, the sets A_i are pairwise disjoint, and $\bigcup_{i=1}^n A_i = \Omega$. Define*

$$\mathcal{G} := \left\{ \bigcup_{i \in I} A_i : I \subseteq \{1, \dots, n\} \right\},$$

the collection of all unions of cells of the partition (including \emptyset and Ω). Then \mathcal{G} is a sigma-algebra on Ω , and it is the smallest sigma-algebra containing every cell A_i . In other words, $\mathcal{G} = \sigma(\mathcal{P})$.

Proof. First, note that $\emptyset \in \mathcal{G}$ (the empty union) and $\Omega \in \mathcal{G}$ (the union of all cells). If $B = \bigcup_{i \in I} A_i$ for some $I \subseteq \{1, \dots, n\}$, then its complement is

$$B^c = \Omega \setminus \bigcup_{i \in I} A_i = \bigcup_{i \notin I} A_i,$$

which again belongs to \mathcal{G} . Hence \mathcal{G} is closed under complements. Since the partition is finite, any countable union of sets in \mathcal{G} reduces to a finite union of distinct cells, and therefore remains in \mathcal{G} . Thus \mathcal{G} is a sigma-algebra.

By construction, each $A_i \in \mathcal{G}$, so \mathcal{G} contains the partition. If \mathcal{F} is any sigma-algebra containing \mathcal{P} , then \mathcal{F} must contain every union of cells (because sigma-algebras are closed under unions). Therefore $\mathcal{G} \subseteq \mathcal{F}$. It follows that \mathcal{G} is the smallest sigma-algebra containing \mathcal{P} , i.e. $\mathcal{G} = \sigma(\mathcal{P})$. \square

Example 1.2 (Coin tosses). Consider two independent tosses of a fair coin with sample space

$$\Omega = \{HH, HT, TH, TT\}.$$

Partition Ω according to the number of heads observed:

$$A_0 = \{TT\}, \quad A_1 = \{HT, TH\}, \quad A_2 = \{HH\}.$$

The sigma-algebra generated by this partition is

$$\sigma(\mathcal{P}) = \{\emptyset, A_0, A_1, A_2, A_0 \cup A_1, A_0 \cup A_2, A_1 \cup A_2, \Omega\}.$$

In economic applications, this construction corresponds to the information available to an agent: observing the partition is the same as knowing only which cell contains the true state, and all measurable events are those that can be distinguished on the basis of this information.

Existence of the sigma-algebra generated by a family

It will be very useful to consider subsigma algebras generated by some events. This is because as an agent it is possible that I have knowledge of some events and I want to know the simplest subsigma algebra that models my knowledge. The next proposition characterises such a situation.

Proposition 1.3. *Let \mathcal{A} be any family of subsets of Ω (or more generally any family of functions $X : \Omega \rightarrow \mathbb{R}$). Then there exists a smallest sigma-algebra containing \mathcal{A} . Concretely, the intersection of all sigma-algebras that contain \mathcal{A} is itself a sigma-algebra and is the smallest sigma-algebra containing \mathcal{A} .*

Proof. Let \mathcal{S} be the collection of all sigma-algebras on Ω that contain \mathcal{A} . Note \mathcal{S} is nonempty because the power set $\mathcal{P}(\Omega)$ is a sigma-algebra containing \mathcal{A} . Define

$$\mathcal{G} := \bigcap_{\mathcal{F} \in \mathcal{S}} \mathcal{F}.$$

Since arbitrary intersections of sigma-algebras are sigma-algebras, \mathcal{G} is a sigma-algebra. By construction \mathcal{G} contains \mathcal{A} , and if \mathcal{H} is any sigma-algebra with $\mathcal{A} \subseteq \mathcal{H}$ then $\mathcal{H} \in \mathcal{S}$, hence $\mathcal{G} \subseteq \mathcal{H}$. Thus \mathcal{G} is the smallest sigma-algebra containing \mathcal{A} . \square

Definition 1.4. If \mathcal{A} is a family of subsets of Ω we write $\sigma(\mathcal{A})$ for the smallest sigma-algebra containing \mathcal{A} ; equivalently $\sigma(\mathcal{A})$ is the intersection of all sigma-algebras that contain \mathcal{A} .

Similarly, if X_1, \dots, X_n are random variables on $(\Omega, \mathcal{F}, \mathbb{P})$, we define

$$\sigma(X_1, \dots, X_n) := \sigma(\{X_1, \dots, X_n\})$$

to be the smallest sigma-algebra with respect to which each X_i is measurable.

Atoms and sigma-algebras generated by random variables

We now give a simple, concrete description of $\sigma(X_1, \dots, X_n)$ in the discrete case. To keep terms explicit we first define an *atom*.

Definition 1.5 (Atom). Let \mathcal{H} be a sigma-algebra on Ω . A nonempty set $A \in \mathcal{H}$ is called an *atom* of \mathcal{H} if for every $E \in \mathcal{H}$ with $E \subseteq A$ we have either $E = \emptyset$ or $E = A$. In other words, an atom is a minimal nonempty event in the sigma-algebra: it cannot be split further inside \mathcal{H} .

Definition 1.6 (Sigma-algebra generated by random variables). Let X_1, \dots, X_n be discrete random variables on $(\Omega, \mathcal{F}, \mathbb{P})$. The sigma-algebra generated by them, denoted $\sigma(X_1, \dots, X_n)$, is the smallest

sigma-algebra that contains all events of the form

$$\{X_i = x\} \quad \text{for each } i = 1, \dots, n \text{ and each value } x \text{ in the range of } X_i.$$

Theorem 1.7. *Let X_1, \dots, X_n be discrete random variables with joint range*

$$\mathcal{S} = \{(X_1(\omega), \dots, X_n(\omega)) : \omega \in \Omega\}.$$

For each $(x_1, \dots, x_n) \in \mathcal{S}$ define the joint-value cell

$$A_{x_1, \dots, x_n} = \{\omega \in \Omega : X_1(\omega) = x_1, \dots, X_n(\omega) = x_n\}.$$

Then the nonempty cells A_{x_1, \dots, x_n} form a partition of Ω , and

$$\sigma(X_1, \dots, X_n) = \sigma(\{A_{x_1, \dots, x_n} : A_{x_1, \dots, x_n} \neq \emptyset\}).$$

Moreover, each nonempty A_{x_1, \dots, x_n} is an atom of $\sigma(X_1, \dots, X_n)$.

Note: I suggest that you see the example below before reading the proof. This proof is straightforward but the notation makes it clunky and complex. The idea is clearly visible in Example 1.8.

Proof. First, for any x_i in the range of X_i , the event $\{X_i = x_i\}$ is a union of certain joint-value cells A_{x_1, \dots, x_n} . Therefore every $\{X_i = x_i\}$ lies in the sigma-algebra generated by the cells, so $\sigma(X_1, \dots, X_n)$ is contained in that sigma-algebra.

Conversely, each joint-value cell A_{x_1, \dots, x_n} can be written as the intersection

$$A_{x_1, \dots, x_n} = \{X_1 = x_1\} \cap \dots \cap \{X_n = x_n\},$$

which shows that every cell belongs to $\sigma(X_1, \dots, X_n)$. Hence the sigma-algebra generated by the cells is contained in $\sigma(X_1, \dots, X_n)$.

The two inclusions together prove equality. Finally, each nonempty A_{x_1, \dots, x_n} is an atom, because no smaller nonempty event in the sigma-algebra can be contained inside it. \square

Example 1.8 (Two coin tosses). Let X and Y be independent coin tosses taking values in $\{0, 1\}$. The possible outcomes are the four pairs

$$(0, 0), \quad (0, 1), \quad (1, 0), \quad (1, 1).$$

Hence the atoms are

$$A_{00} = \{\omega : X = 0, Y = 0\}, \quad A_{01} = \{\omega : X = 0, Y = 1\},$$

$$A_{10} = \{\omega : X = 1, Y = 0\}, \quad A_{11} = \{\omega : X = 1, Y = 1\}.$$

Therefore

$$\sigma(X, Y) = \{\emptyset, A_{00}, A_{01}, A_{10}, A_{11}, A_{00} \cup A_{01}, \dots, \Omega\},$$

i.e. all unions of the four atoms. In general, $\sigma(X, Y)$ corresponds to knowing the joint outcome of (X, Y) .

This construction shows how multiple discrete random variables induce a partition of the sample space into joint-value cells, and the generated sigma-algebra is precisely the information structure that distinguishes which cell the true state belongs to.

1.4 Measurable Functions

Let (Ω, \mathcal{F}) be a measurable space, where \mathcal{F} is a sigma-algebra on Ω . A function $X : \Omega \rightarrow \mathbb{R}$ is said to be measurable with respect to \mathcal{F} if for every real number b , the set

$$E_b = \{\omega \in \Omega : X(\omega) \leq b\}$$

belongs to \mathcal{F} . In other words, the event that X takes a value less than or equal to any given threshold must be describable in terms of the information encoded by \mathcal{F} .

When \mathcal{F} is generated by a finite partition $\mathcal{P} = \{A_1, \dots, A_n\}$, the structure of measurable functions becomes very transparent. For each cell A_i we define the indicator $\mathbf{1}_{A_i}$, which is 1 on A_i and 0 elsewhere.

Theorem 1.9. *Let Ω be a set and let $\mathcal{P} = \{A_1, \dots, A_n\}$ be a finite partition of Ω . Write $\mathcal{F} = \sigma(\mathcal{P})$ for the sigma-algebra generated by the partition. If $X : (\Omega, \mathcal{F}) \rightarrow \mathbb{R}$ is \mathcal{F} -measurable (in the sense that for every real b the set*

$$E_b = \{\omega \in \Omega : X(\omega) \leq b\}$$

lies in \mathcal{F}), then X is constant on each cell A_i . Consequently there exist real numbers c_1, \dots, c_n such that

$$X = \sum_{i=1}^n c_i \mathbf{1}_{A_i}.$$

Moreover, this representation is unique: c_i is the common value of X on A_i .

Proof. Because $\mathcal{F} = \sigma(\mathcal{P})$, every set in \mathcal{F} is a union of cells A_i . In particular each set E_b is a union of some subcollection of the cells.

Fix an index i and suppose, for contradiction, that X is not constant on A_i . Then there exist $\omega_1, \omega_2 \in A_i$ with $X(\omega_1) \neq X(\omega_2)$. Without loss of generality assume $X(\omega_1) < X(\omega_2)$. Choose a real number b with

$$X(\omega_1) < b < X(\omega_2).$$

By definition $\omega_1 \in E_b$ and $\omega_2 \notin E_b$. But $E_b \in \mathcal{F}$, so E_b must be a union of whole cells of the partition; it cannot split a single cell into two parts. This contradicts the fact that ω_1 and ω_2 lie

in the same cell A_i while one belongs to E_b and the other does not. Therefore no such pair ω_1, ω_2 exists, and X must be constant on A_i .

Since the choice of i was arbitrary, X is constant on every cell. Let c_i denote the common value of X on A_i . Then for every $\omega \in \Omega$ we have $X(\omega) = \sum_{i=1}^n c_i \mathbf{1}_{A_i}(\omega)$, which gives the claimed representation. Uniqueness of the coefficients c_i follows because the A_i are disjoint: the value of X on A_i determines c_i . \square

In summary, any \mathcal{F} -measurable function X must be constant on each A_i , since no finer distinction is possible within the sigma-algebra and every measurable function has the representation

$$X = \sum_{i=1}^n c_i \mathbf{1}_{A_i},$$

where $c_i \in \mathbb{R}$ is the constant value of X on the set A_i .

This characterization shows concretely that measurability with respect to a partition means depending only on which cell contains the true state of the world. Later, this perspective will serve as the foundation for conditional expectation, where general random variables are approximated by measurable ones relative to a chosen sigma-algebra.

Remark 1.10. All of the results above extend immediately from finite to countable partitions. If $\mathcal{P} = \{A_1, A_2, \dots\}$ is a countable partition of Ω , then the sigma-algebra generated by \mathcal{P} consists of all unions of the cells A_i , which may now be countably infinite. Every $\sigma(\mathcal{P})$ -measurable function is constant on each A_i , and therefore can be written in the form

$$X = \sum_{i=1}^{\infty} c_i \mathbf{1}_{A_i},$$

where the series is pointwise finite because each $\omega \in \Omega$ belongs to exactly one cell. Thus measurable functions with respect to countable partitions are exactly the discrete random variables: they take only countably many distinct values, one for each cell. In this sense, our discussion provides a complete account of discrete random variables.

2 Probability Measures and Expectation

2.1 Probability Measure

Let (Ω, \mathcal{F}) be a measurable space. A *probability measure* is a function $P : \mathcal{F} \rightarrow [0, 1]$ satisfying the following axioms:

1. $P(\Omega) = 1$,

2. P is countably additive: if $(A_i)_{i=1}^{\infty}$ is a sequence of disjoint sets in \mathcal{F} , then

$$P\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} P(A_i).$$

The triple (Ω, \mathcal{F}, P) is called a *probability space*.

2.2 Expectation

Let $X : (\Omega, \mathcal{F}) \rightarrow \mathbb{R}$ be a random variable. Its expectation is denoted by the “integral notation”

$$\mathbb{E}[X] = \int_{\Omega} X(\omega) dP(\omega).$$

This definition is completely general and unifies the discrete and continuous cases.

In the discrete case, where Ω is finite or countable and P is determined by point masses $P(\{\omega\})$, the expectation reduces to a weighted sum:

$$\mathbb{E}[X] = \sum_{\omega \in \Omega} X(\omega) P(\{\omega\}).$$

In the continuous case, where $\Omega \subseteq \mathbb{R}$ and P admits a density f , the expectation takes the familiar form

$$\mathbb{E}[X] = \int_{\Omega} X(x) f(x) dx.$$

2.3 Example

Let $\Omega = [0, 1]$, \mathcal{F} the usual sigma-algebra generated by intervals, and P be the Lebesgue measure restricted to $[0, 1]$. This is the uniform distribution on $[0, 1]$. For a random variable $X(\omega) = \omega$, the expectation is

$$\mathbb{E}[X] = \int_0^1 \omega d\omega = \frac{1}{2}.$$

Thus under the uniform distribution the mean value is $1/2$, as expected.

2.4 Discussion

The material developed so far can be viewed as a complete account of the theory of discrete random variables. Sigma-algebras generated by finite or countable partitions describe the information structure, and measurable functions with respect to these sigma-algebras are precisely the discrete random variables. Each such random variable is constant on every cell of the partition and can be written as a (finite or countable) linear combination of indicator functions.

2.5 Exercises

Exercise 1. Let $\Omega = \{1, 2, 3, 4\}$. For each of the following partitions \mathcal{P} of Ω , explicitly write down the sigma-algebra $\sigma(\mathcal{P})$ as a collection of subsets of Ω .

1. $\mathcal{P} = \{\{1, 2\}, \{3, 4\}\}$.
2. $\mathcal{P} = \{\{1\}, \{2\}, \{3, 4\}\}$.
3. $\mathcal{P} = \{\{1\}, \{2, 3\}, \{4\}\}$.

Explain in each case why your collection is closed under complements and countable unions, and why no smaller sigma-algebra contains the partition.

Exercise 2. Consider the sample space $\Omega = \{HH, HT, TH, TT\}$ of two coin tosses. Let $\mathcal{P} = \{A_0, A_1, A_2\}$ be the partition according to the number of heads:

$$A_0 = \{TT\}, \quad A_1 = \{HT, TH\}, \quad A_2 = \{HH\}.$$

Define a random variable $X : \Omega \rightarrow \mathbb{R}$ by $X(\omega) = \text{number of heads in outcome } \omega$. Show that X is measurable with respect to $\sigma(\mathcal{P})$ and write it in the form

$$X = \sum_{i=0}^2 c_i \mathbf{1}_{A_i}.$$

Is the random variable $Y(\omega) = \mathbf{1}_{\{\omega=HT\}}$ measurable with respect to $\sigma(\mathcal{P})$? Explain.

3 Markov Decision Processes

3.1 Motivation

Economic agents frequently face decisions under uncertainty and over time. A worker chooses whether to accept or reject a job offer today, knowing that the consequences of this choice affect both present and future payoffs. A firm decides whether to invest in capital this period, taking into account the stochastic evolution of productivity tomorrow. These problems combine three elements:

- (i) an evolving random environment,
- (ii) a sequence of decisions that must be taken contingent on the realized states, and
- (iii) an objective function, typically the expected present discounted value of returns.

The natural mathematical framework for such problems is a *Markov Decision Process (MDP)*. An MDP formalizes the primitives of the problem: a state space describing the relevant economic environment, an action space describing feasible choices, a stochastic law of motion for states, and a reward or payoff function. A policy specifies how actions are chosen as a function of states, and the central object of interest is the *value function*, defined as the expected discounted payoff when starting from a given state and following a particular policy.

The recursive formulation of dynamic programming is essential here. Rather than treating the problem as one of optimizing over an infinite sequence of actions, we exploit the *principle of optimality*: any optimal plan, when restricted to a subproblem beginning at a future state, must itself be optimal for that subproblem. This principle yields the *Bellman equation*, a recursive characterization of the value function that collapses an infinite-horizon optimization problem into a fixed-point problem.

To make the framework concrete, we turn to the McCall job search problem. Here an unemployed worker receives wage offers drawn from a known distribution. Each period she chooses whether to accept the offer and exit unemployment, or reject it and continue searching. The model illustrates all the primitives of an MDP: the state is the current offer, the action is to accept or reject, the payoff depends on the action, and the law of motion is given by the stochastic distribution of future offers. The Bellman equation captures the trade-off between immediate acceptance and the option value of continued search. This simple example will serve as our running case study.

3.2 Definitions

We work in a probability space where all randomness is generated by a sequence of random variables. At each date $t = 0, 1, 2, \dots$, the system is in a state S_t taking values in a finite set S . The agent chooses an action A_t from a finite set A . The choice of action affects both the immediate reward and the distribution of the next state.

Formally, suppose that for every pair $(s, a) \in S \times A$ we are given two primitives. First, the one-period reward $r(s, a)$ is the payoff the agent receives if she takes action a in state s . Second, the transition law is described by conditional probabilities

$$\mathbb{P}(S_{t+1} = s' \mid S_t = s, A_t = a) = P(s' \mid s, a),$$

which specify the distribution of the next state random variable given the current state and action. Together, $\{r(s, a), P(\cdot \mid s, a)\}$ describe the environment.

A stationary policy π is a function from states to actions. It specifies how the agent acts: whenever the current state is s , the chosen action is $A_t = \pi(s)$. Thus the pair (S_t, A_t) evolves as a controlled Markov chain. More general policies could be random or history-dependent, but we restrict attention to stationary deterministic policies, since under discounting they are sufficient for optimality in the finite case.

Given an initial state $S_0 = s$, a policy π , and a discount factor $\beta \in (0, 1)$, the value function is

the expected present discounted value of rewards,

$$V^\pi(s) := \mathbb{E}_s^\pi \left[\sum_{t=0}^{\infty} \beta^t r(S_t, \pi(S_t)) \right],$$

where the expectation is taken with respect to the probability law generated by π and the transition probabilities. This value function records the performance of policy π when starting in state s .

The optimal value function is defined by

$$V^*(s) := \sup_{\pi} V^\pi(s),$$

which gives the maximal attainable expected payoff from state s . An optimal policy is one that achieves this value.

The recursive characterization of V^* is given by the Bellman equation. For each state s ,

$$V^*(s) = \max_{a \in A} \left\{ r(s, a) + \beta \mathbb{E}[V^*(S_{t+1}) \mid S_t = s, A_t = a] \right\}.$$

The expectation on the right is computed using the transition probabilities $P(s' \mid s, a)$. Thus, in the finite case,

$$V^*(s) = \max_{a \in A} \left\{ r(s, a) + \beta \sum_{s' \in S} P(s' \mid s, a) V^*(s') \right\}.$$

This is the formula we will actually use. It states that the optimal value of starting in state s is the maximum, over feasible actions a , of the immediate reward plus the discounted expected continuation value.

3.3 Examples: The McCall Job Search Problem

Consider an unemployed worker who at each date $t = 0, 1, 2, \dots$ receives a wage offer W_t . The sequence $\{W_t\}$ is i.i.d. with cumulative distribution function F on \mathbb{R}_+ . If the worker accepts the offer $W_t = w$, she exits unemployment permanently and earns w each period thereafter. If she rejects, she receives an unemployment benefit $b \geq 0$ and faces a new offer W_{t+1} next period. The discount factor is $\beta \in (0, 1)$.

The state is the current offer w , and the actions are *accept* or *reject*. The payoff from acceptance is w each period, yielding a present value of $w/(1 - \beta)$. If the worker rejects, she gets b today and expects to continue tomorrow with a new offer. The value function satisfies the Bellman equation

$$V(w) = \max \left\{ \frac{w}{1 - \beta}, b + \beta \mathbb{E}[V(W')] \right\},$$

where $W' \sim F$ is an independent wage draw.

Reservation wage: The structure of the Bellman equation implies there exists a threshold w^* such that the worker accepts if $w \geq w^*$ and rejects otherwise. At the reservation wage the worker

is indifferent, so w^* satisfies

$$\frac{w^*}{1 - \beta} = b + \beta \mathbb{E}[V(W')].$$

Notice that if the worker rejects, the continuation value does not depend on today's offer. Thus we can write

$$\mathbb{E}[V(W')] = \int_0^{w^*} V(w^*) dF(w') + \int_{w^*}^{\infty} \frac{w'}{1 - \beta} dF(w').$$

The first term reflects the event that tomorrow's offer is below w^* , in which case the worker will reject again and obtain continuation value $V(w^*)$. The second term covers the case that tomorrow's offer exceeds w^* , in which case she accepts at w' . Since $V(w^*) = w^*/(1 - \beta)$, we obtain the fundamental equation for the reservation wage:

$$\frac{w^*}{1 - \beta} = b + \beta \left[F(w^*) \cdot \frac{w^*}{1 - \beta} + \int_{w^*}^{\infty} \frac{w'}{1 - \beta} dF(w') \right].$$

Multiplying through by $(1 - \beta)$ gives the usable formula

$$w^* = (1 - \beta)b + \beta \left(F(w^*) w^* + \int_{w^*}^{\infty} w' dF(w') \right).$$

This nonlinear equation in w^* can be solved numerically given b, β , and F . In the discrete case the integrals become finite sums.

Discussion: The reservation wage increases with the benefit level b : higher benefits raise the value of waiting, leading to more selectivity. It also increases with the discount factor β , since patience magnifies the value of future opportunities. The shape of the wage distribution matters: if high wages are sufficiently likely, the integral term is larger, raising w^* . Conversely, if the distribution is concentrated at low wages, w^* is lower.

Economic interpretations: In countries with generous unemployment insurance and high expected wages, workers set higher reservation wages, which lengthens unemployment spells but leads to better job matches. In countries with minimal benefits and lower wage distributions, workers accept more quickly, reducing unemployment duration but often at the cost of poorer matches and lower long-term income. Thus, differences in labor market institutions and productivity distributions translate directly into observable differences in unemployment and standards of living.

3.4 Discussion

The concepts of policy and value function admit a direct economic interpretation. A policy π is nothing more than a rule of behavior, prescribing which action is taken as a function of the current state. In economic models it plays the role of a decision rule: in the McCall problem it is the acceptance or rejection rule; in investment models it is the mapping from the capital stock to the investment decision. The value function V^π represents the performance of such a rule. It answers the question: if the agent follows policy π from now on, what is the expected present discounted

value of her payoffs starting from a given state? In this sense, the value function is the evaluation of a particular behavioral plan.

The optimal value function V^* is then the upper envelope of all feasible policies. It embodies the idea that rational economic agents compare policies, discard inferior ones, and act as if maximizing over the feasible set. An optimal policy π^* realizes this maximum. Thus the recursive framework provides a precise link between rules of behavior and their economic consequences, without requiring the agent to optimize over entire sequences of actions all at once.

The recursive formulation, captured in the Bellman equation, is at the heart of Sargent’s framework for macroeconomics. Instead of treating a dynamic optimization problem as an intertemporal choice over infinite paths, the recursive approach reduces it to the study of a functional equation for V^* . The principle of optimality guarantees that once the Bellman equation is solved, the associated policy is globally optimal. This transformation has two decisive advantages. First, it delivers a transparent economic interpretation: each decision balances current reward against the discounted expected continuation value, with the latter summarizing all future possibilities. Second, it makes computation feasible. The contraction mapping property ensures that repeated application of the Bellman operator converges rapidly to V^* , which is the basis of modern numerical dynamic programming.

In Sargent’s recursive macroeconomics, this framework is not a technical aside but the central organizing device. Models of consumption, saving, investment, and search are all formulated as recursive problems, with state variables capturing the relevant information and Bellman equations delivering both the qualitative logic and the quantitative solution. The job search model of McCall is the simplest laboratory: the state is a single random wage offer, the policy is a reservation wage rule, and the Bellman equation encodes the trade-off between immediate employment and the option value of waiting. More elaborate models—capital accumulation, stochastic growth, asset pricing—are built on precisely the same recursive foundations.

3.5 Exercises

Exercise 3.1. Compute the reservation wage and the corresponding value in the McCall model when wage offers are uniformly distributed on $[0, M]$. Let the unemployment benefit be $b \geq 0$, the discount factor be $\beta \in (0, 1)$, and assume the reservation wage lies in $[0, M]$. Derive the reservation wage w^* and the value at the reservation wage.

Final answer:

$$w^* = \frac{M - \sqrt{(1 - \beta) M (M(1 + \beta) - 2\beta b)}}{\beta}, \quad V(w^*) = \frac{w^*}{1 - \beta}.$$

Exercise 3.2. Solve a small finite-state macroeconomic MDP by matrix methods. Consider an economy whose productivity can be either high (H) or low (L). Each period a planner (or a representative firm) chooses between two actions: “invest” (I) or “do nothing” (N). If the economy

is in state H and the planner chooses I the immediate payoff is $r(H, I) = 10$; if in H and she chooses N the payoff is $r(H, N) = 8$. If the state is L the payoffs are $r(L, I) = 6$ and $r(L, N) = 3$. Transition probabilities depend on the action: under I the transition matrix (rows current state H, L ; columns next state H, L) is

$$P(\cdot | \cdot, I) = \begin{pmatrix} 0.9 & 0.1 \\ 0.6 & 0.4 \end{pmatrix},$$

and under N it is

$$P(\cdot | \cdot, N) = \begin{pmatrix} 0.7 & 0.3 \\ 0.3 & 0.7 \end{pmatrix}.$$

Take $\beta = 0.95$. State the MDP formally in words, construct the four deterministic stationary policies, and compute each policy's value vector by solving the linear system. Explain how to obtain the optimal value function and an optimal policy from these computations.

Exercise 3.3. Derive the general linear-algebra principle used to solve finite-state MDPs and give the principal constructive corollaries used in applications. (Answer in one paragraph.)

Answer: For a finite-state MDP with n states and a fixed stationary deterministic policy π , policy evaluation reduces to a linear system: if P^π is the $n \times n$ transition matrix induced by π and $r^\pi \in \mathbb{R}^n$ is the vector of one-period rewards under π , then the value vector satisfies

$$V^\pi = r^\pi + \beta P^\pi V^\pi,$$

so that, because $\beta \in (0, 1)$, the matrix $I - \beta P^\pi$ is invertible and

$$V^\pi = (I - \beta P^\pi)^{-1} r^\pi.$$

4 Conditional Expectation

4.1 Motivation

In economics, uncertainty is rarely just about which state of the world occurs; it is equally about what agents know when they make decisions. The central tool for formalizing this interplay between randomness and information is *conditional expectation*.

From a statistical perspective, conditioning means updating beliefs when new information arrives. In introductory courses this is presented as conditioning on an event, for example learning that a die roll landed in the set $\{1, 2, 3\}$. In more advanced applications, especially in dynamic models, information typically comes in the form of a sigma-algebra: a collection of events that captures precisely what an agent can distinguish at a given point in time. Conditional expectation with respect to such a sigma-algebra is the mathematical device that transforms a random variable into its best forecast based solely on that information.

For economists, this idea is indispensable. In dynamic programming, the Bellman equation rests on taking expected future payoffs conditional on the current state. In financial economics, pricing kernels and asset returns are defined through conditional expectations given market information. In econometrics, forecasting and estimation rely on projecting random variables onto the space of what is observable. Across all of these contexts, conditional expectation provides the rigorous link between information and decision making, ensuring that our models respect both the randomness of outcomes and the informational constraints faced by agents.

In macroeconomics, the role of conditional expectation is particularly visible in the debate between the adaptive expectations hypothesis (AEH) and the rational expectations hypothesis (REH). Under AEH, agents form forecasts of future variables by adjusting past forecasts in response to recent forecast errors; the updating rule is ad hoc and does not directly appeal to probability theory. By contrast, REH formalizes expectations as mathematical conditional expectations: given the information available at time t , agents' forecasts of tomorrow's variables are defined as the conditional expectation with respect to that sigma-algebra. This shift is profound. It means that expectations are model-consistent, forward-looking, and disciplined by probability theory rather than by arbitrary adjustment rules. In this way, conditional expectation serves as the mathematical backbone of modern macroeconomic modeling, ensuring that the way agents form beliefs is internally consistent with the structure of the economy itself.

4.2 Conditional Expectation with respect to Events

Theorem 4.1. *Let (Ω, \mathcal{F}, P) be a probability space and let $A \in \mathcal{F}$ satisfy $P(A) > 0$. For any integrable random variable $X : \Omega \rightarrow \mathbb{R}$,*

$$\mathbb{E}[X|A] := \int_{\Omega} X dP_A = \frac{1}{P(A)} \int_{\Omega} X \mathbf{1}_A dP$$

where $P_A(B) = P(B | A) = P(B \cap A)/P(A)$.

Proof. We will prove the claim only for Ω finite or countable we may write expectations as sums. For each $\omega \in \Omega$ let $p(\omega) = P(\{\omega\})$. The conditional probability mass function under P_A is

$$p_A(\omega) = P_A(\{\omega\}) = \frac{P(\{\omega\} \cap A)}{P(A)} = \begin{cases} \frac{p(\omega)}{P(A)}, & \omega \in A, \\ 0, & \omega \notin A. \end{cases}$$

Hence the expectation of X under P_A is the sum

$$\mathbb{E}_{P_A}[X] = \sum_{\omega \in \Omega} X(\omega) p_A(\omega) = \sum_{\omega \in A} X(\omega) \frac{p(\omega)}{P(A)} = \frac{1}{P(A)} \sum_{\omega \in \Omega} X(\omega) \mathbf{1}_A(\omega) p(\omega).$$

The final sum is precisely $\frac{1}{P(A)} \int_{\Omega} X \mathbf{1}_A dP$, so

$$\mathbb{E}[X|A] = \frac{1}{P(A)} \int_{\Omega} X \mathbf{1}_A dP,$$

which proves the claim. \square

Note that in the continuous case, when P admits a density f with respect to Lebesgue measure, one has

$$\mathbb{E}[X | A] = \frac{\int_A X(\omega) f(\omega) d\omega}{\int_A f(\omega) d\omega}.$$

Example 4.2. Let $\Omega = \{H, T\}^3$ with the uniform probability $P(\{\omega\}) = 1/8$. Let $X(\omega)$ be the number of heads in outcome ω . For an atom A of a sigma-algebra the formula from the theorem gives

$$\mathbb{E}[X | A] = \frac{\sum_{\omega \in A} X(\omega) P(\{\omega\})}{P(A)} = \frac{\sum_{\omega \in A} X(\omega) (1/8)}{P(A)}.$$

Conditioning on \mathcal{F}_1 . The atoms are

$$A_H = \{HHH, HHT, HTH, HTT\}, \quad A_T = \{THH, THT, TTH, TTT\},$$

each has probability $P(A_H) = P(A_T) = 4/8 = 1/2$. Compute the numerator sums:

$$\sum_{\omega \in A_H} X(\omega) = 3 + 2 + 2 + 1 = 8, \quad \sum_{\omega \in A_T} X(\omega) = 2 + 1 + 1 + 0 = 4.$$

Hence

$$\begin{aligned} \mathbb{E}[X | A_H] &= \frac{(1/8) \cdot 8}{1/2} = \frac{1}{1/2} = 2, \\ \mathbb{E}[X | A_T] &= \frac{(1/8) \cdot 4}{1/2} = \frac{1/2}{1/2} = 1. \end{aligned}$$

Thus $\mathbb{E}[X | \mathcal{F}_1] = 2$ on A_H and 1 on A_T .

Conditioning on \mathcal{F}_2 . The atoms are the four outcomes of the first two tosses, each containing

two full outcomes and each with probability $2/8 = 1/4$. Their numerator sums are

$$\begin{aligned}\sum_{\omega: \text{first two} = HH} X(\omega) &= 3 + 2 = 5, \\ \sum_{\omega: \text{first two} = HT} X(\omega) &= 2 + 1 = 3, \\ \sum_{\omega: \text{first two} = TH} X(\omega) &= 2 + 1 = 3, \\ \sum_{\omega: \text{first two} = TT} X(\omega) &= 1 + 0 = 1.\end{aligned}$$

Hence

$$\begin{aligned}\mathbb{E}[X \mid \text{first two} = HH] &= \frac{(1/8) \cdot 5}{1/4} = \frac{5/8}{1/4} = 2.5, \\ \mathbb{E}[X \mid \text{first two} = HT] &= \frac{(1/8) \cdot 3}{1/4} = \frac{3/8}{1/4} = 1.5, \\ \mathbb{E}[X \mid \text{first two} = TH] &= 1.5, \\ \mathbb{E}[X \mid \text{first two} = TT] &= \frac{(1/8) \cdot 1}{1/4} = \frac{1/8}{1/4} = 0.5.\end{aligned}$$

Equivalently, if $k \in \{0, 1, 2\}$ denotes the number of heads in the first two tosses then

$$\mathbb{E}[X \mid \mathcal{F}_2] = k + \frac{1}{2} \quad \text{on the atom with } k \text{ heads.}$$

Conditioning on \mathcal{F}_3 . Each atom is a singleton $\{\omega\}$ with $P(\{\omega\}) = 1/8$, so

$$\mathbb{E}[X \mid \{\omega\}] = \frac{(1/8)X(\omega)}{1/8} = X(\omega).$$

Hence $\mathbb{E}[X \mid \mathcal{F}_3] = X$.

These calculations follow directly from the event-conditioning formula $\mathbb{E}[X \mid A] = \frac{\mathbb{E}[X\mathbf{1}_A]}{P(A)}$ applied to each atom A of the relevant sigma-algebra.

4.3 Conditional expectation via partitions: an economic perspective

Consider an agent who, before any information is revealed, has a belief about a random variable X . Her ex ante belief is the unconditional expectation $\mathbb{E}[X]$. Now suppose a natural experiment occurs that reveals only coarse information: the outcome ω of the underlying experiment is not observed directly, but only through which cell of a partition $\{A_1, \dots, A_n\}$ it belongs to. Each cell A_i represents an information set. The question is: how should the agent update her expectation of X after seeing which cell occurs?

Definition 4.3 (Conditional expectation with respect to a partition sigma algebra). Let (Ω, \mathcal{F}, P)

be a probability space, and let $\{A_1, \dots\}$ be a partition of Ω . Let $\mathcal{G} = \sigma(\{A_1, \dots\})$ be the sigma-algebra generated by this partition. For any random variable $X : \Omega \rightarrow \mathbb{R}$ we define

$$\mathbb{E}[X \mid \mathcal{G}](\omega) = \sum_{i=1}^{\infty} \mathbb{E}[X \mid A_i] \mathbf{1}_{A_i}(\omega),$$

where

$$\mathbb{E}[X \mid A_i] = \frac{\int_{A_i} X dP}{P(A_i)} \quad \text{for } P(A_i) > 0.$$

Economic interpretation. Ex ante the agent's expectation is $\mathbb{E}[X]$. Ex post, after learning that the realized state lies in some cell A_i , she updates to the conditional expectation $\mathbb{E}[X \mid A_i] = m_i$. The conditional expectation $\mathbb{E}[X \mid \mathcal{G}]$ therefore bundles together all these ex post expectations into a single random variable: it tells us, for each possible outcome, which updated forecast the agent would hold given the information that outcome reveals.

4.4 Conditional Expectation: Defining property

Now we prove an important property that defines conditional expectation in a more general setting.

Theorem 4.4 (Existence and uniqueness of conditional expectation). *Let (Ω, \mathcal{F}, P) be a probability space and let $\mathcal{G} \subseteq \mathcal{F}$ be a sub-sigma-algebra. If $X : \Omega \rightarrow \mathbb{R}$ is integrable (i.e. $\int_{\Omega} |X| dP < \infty$), then there exists a \mathcal{G} -measurable integrable random variable Y such that*

$$\int_A Y dP = \int_A X dP \quad \text{for every } A \in \mathcal{G}.$$

Any two such \mathcal{G} -measurable functions agree almost surely (i.e. they are equal except on a set of probability zero). A random variable Y with the displayed property is denoted $\mathbb{E}[X \mid \mathcal{G}]$ and is called a conditional expectation of X given \mathcal{G} .

Proof of existence and uniqueness in the partition case. We prove the theorem in the special case when $\mathcal{G} = \sigma(\mathcal{P})$ is the sigma-algebra generated by a finite or countable partition $\mathcal{P} = \{A_i : i \in I\}$ of Ω . The proof is constructive and elementary, and it both exhibits existence and proves uniqueness (up to null sets).

Construction / Existence. For each index $i \in I$ with $P(A_i) > 0$ define

$$m_i := \frac{\int_{A_i} X dP}{P(A_i)},$$

the average (expectation) of X on the atom A_i . If $P(A_i) = 0$ define $m_i := 0$. Now set

$$Y(\omega) := \sum_{i \in I} m_i \mathbf{1}_{A_i}(\omega).$$

By construction Y is \mathcal{G} -measurable (it is constant on each part) and integrable (since $\sum_i |m_i|P(A_i) \leq \sum_i \int_{A_i} |X| dP = \int_{\Omega} |X| dP < \infty$).

Let $A \in \mathcal{G}$. Then A is a disjoint union of parts: $A = \bigcup_{j \in J} A_{i_j}$ for some index set $J \subseteq I$. Using linearity of the integral and the definition of m_i ,

$$\int_A Y dP = \sum_{j \in J} m_{i_j} P(A_{i_j}) = \sum_{j \in J} \int_{A_{i_j}} X dP = \int_A X dP.$$

Thus Y satisfies the defining property of conditional expectation on every $A \in \mathcal{G}$. This proves existence in the partition case.

Uniqueness (up to a null set). Suppose Y' is another \mathcal{G} -measurable integrable function with $\int_A Y' dP = \int_A X dP$ for every $A \in \mathcal{G}$. Then for any atom A_i with $P(A_i) > 0$ taking $A = A_i$ gives

$$\int_{A_i} (Y - Y') dP = 0.$$

But both Y and Y' are constant on A_i , so $Y - Y'$ is constant on A_i . Therefore the integral of $Y - Y'$ over A_i equals that constant times $P(A_i)$; since $P(A_i) > 0$ this forces the constant to be zero, i.e. $Y = Y'$ on A_i . Since this holds for every part with positive probability, $Y = Y'$ almost surely. Hence the conditional expectation is unique up to sets of probability zero.

This completes the constructive proof of existence and uniqueness when \mathcal{G} is the sigma-algebra generated by a partition. The general theorem (for arbitrary sub-sigma-algebras) is true as stated above; its proof in full generality requires additional measure-theoretic machinery (Radon–Nikodym theorem or standard Hilbert-space/approximation arguments) which we omit here and replace in applications by the explicit partition construction that is all we need for discrete models. \square

Remark 4.5 (Econometric interpretation). In econometrics, conditional expectation appears naturally in the theory of forecasting. Suppose an econometrician wishes to predict a random variable Y using information contained in a sigma-algebra \mathcal{G} , which represents the observations or regressors that have been recorded. A forecast Y^e is called optimal if it minimizes the mean squared error

$$\mathbb{E}[(Y - Y^e)^2]$$

among all \mathcal{G} -measurable candidates Y^e .

The first-order condition for this minimization states that the forecast error $Y - Y^e$ is orthogonal to every \mathcal{G} -measurable variable. In particular, for every $A \in \mathcal{G}$ we have

$$\mathbb{E}[(Y - Y^e) \mathbf{1}_A] = 0.$$

Rearranging yields

$$\mathbb{E}[Y \mathbf{1}_A] = \mathbb{E}[Y^e \mathbf{1}_A],$$

which is exactly the defining property of conditional expectation from the theorem above and thus

$$Y^e = \mathbb{E}[X \mid \mathcal{G}].$$

Thus, in econometric language, conditional expectation is nothing more than the best forecast of Y based on the available information \mathcal{G} , in the sense of minimizing mean squared error.

4.5 Conditional Expectation with respect to a Random Variable

It is often useful to speak of the conditional expectation of one random variable given another. Formally, if X and Y are random variables on (Ω, \mathcal{F}, P) , we define

$$\mathbb{E}[X \mid Y] := \mathbb{E}[X \mid \sigma(Y)],$$

where $\sigma(Y)$ denotes the sigma-algebra generated by Y . That is, $\mathbb{E}[X \mid Y]$ is the conditional expectation of X given the information contained in Y . By construction, $\mathbb{E}[X \mid Y]$ is a random variable that is measurable with respect to $\sigma(Y)$, and therefore can be expressed as a function of Y ¹:

$$\mathbb{E}[X \mid Y](\omega) = g(Y(\omega)) \quad \text{for some function } g : \mathbb{R} \rightarrow \mathbb{R}.$$

Examples.

1. *Discrete case.* Let Y take values y_1, \dots, y_n with $P(Y = y_j) > 0$. Then

$$\mathbb{E}[X \mid Y = y_j] = \frac{\int_{\{Y=y_j\}} X dP}{P(Y = y_j)},$$

and $\mathbb{E}[X \mid Y]$ is the random variable that equals this conditional mean on the event $\{Y = y_j\}$:

$$\mathbb{E}[X \mid Y](\omega) = \sum_{j=1}^n \mathbb{E}[X \mid Y = y_j] \mathbf{1}_{\{Y=y_j\}}(\omega).$$

2. *Coin toss illustration.* Toss two fair coins. Let Y be the first coin (1 if H , 0 if T), and let X be the total number of heads. Then $\mathbb{E}[X \mid Y = 1] = 1.5$ (one head already, plus 1/2 expected from the second coin), while $\mathbb{E}[X \mid Y = 0] = 0.5$. Hence

$$\mathbb{E}[X \mid Y](\omega) = \begin{cases} 1.5 & \text{if the first toss is } H, \\ 0.5 & \text{if the first toss is } T. \end{cases}$$

¹In general, this follows from Doob-Dynkin lemma which we will not prove. For discrete Y , it follows from the characterisations in Chapter 2.

3. *Continuous case.* If (X, Y) has joint density $f_{X,Y}$ and Y has marginal density f_Y , then

$$\mathbb{E}[X \mid Y = y] = \frac{\int_{\mathbb{R}} x f_{X,Y}(x, y) dx}{f_Y(y)}.$$

The conditional expectation $\mathbb{E}[X \mid Y]$ is the function of Y obtained by substituting the realized value $Y(\omega)$ into this formula.

4.6 Properties of conditional expectation

Let $\mathcal{P} = \{A_1, A_2, \dots\}$ be a finite or countable partition of Ω , and let

$$\mathcal{G} = \sigma(\mathcal{P})$$

be the sigma-algebra it generates. Recall that for any random variable X ,

$$\mathbb{E}[X \mid \mathcal{G}](\omega) = \sum_i \frac{\int_{A_i} X dP}{P(A_i)} \mathbf{1}_{A_i}(\omega).$$

Theorem 4.6 (Linearity). *For random variables X, Z and constants $a, b \in \mathbb{R}$,*

$$\mathbb{E}[aX + bZ \mid \mathcal{G}] = a\mathbb{E}[X \mid \mathcal{G}] + b\mathbb{E}[Z \mid \mathcal{G}].$$

Proof. It suffices to check the claim on each part A_i ,

$$\mathbb{E}[aX + bZ \mid A_i] = \frac{\int_{A_i} (aX + bZ) dP}{P(A_i)} = a\mathbb{E}[X \mid A_i] + b\mathbb{E}[Z \mid A_i].$$

Hence the conditional expectation is linear. □

Proof. By the partition formula, for $\mathcal{G} = \sigma(\mathcal{P})$ we have

$$\mathbb{E}[X \mid \mathcal{G}] = \sum_i \frac{\int_{A_i} X dP}{P(A_i)} \mathbf{1}_{A_i}.$$

Now condition again on $\mathcal{H} = \sigma(\mathcal{Q})$, where each $B_j \in \mathcal{Q}$ is a union of parts $A_i \in \mathcal{P}$. Then

$$\mathbb{E}[\mathbb{E}[X \mid \mathcal{G}] \mid \mathcal{H}] = \sum_j \frac{\int_{B_j} \mathbb{E}[X \mid \mathcal{G}] dP}{P(B_j)} \mathbf{1}_{B_j}.$$

Substitute the formula for $\mathbb{E}[X \mid \mathcal{G}]$:

$$\int_{B_j} \mathbb{E}[X \mid \mathcal{G}] dP = \sum_{i: A_i \subseteq B_j} \frac{\int_{A_i} X dP}{P(A_i)} P(A_i) = \sum_{i: A_i \subseteq B_j} \int_{A_i} X dP = \int_{B_j} X dP.$$

Hence

$$\mathbb{E}[\mathbb{E}[X | \mathcal{G}] | \mathcal{H}] = \sum_j \frac{\int_{B_j} X dP}{P(B_j)} \mathbf{1}_{B_j} = \mathbb{E}[X | \mathcal{H}].$$

□

Remark 4.7 (Econometric interpretation of the tower property). The tower property has a simple forecasting interpretation. Suppose there are two econometricians: one has access to the finer information set \mathcal{G} , while another has access only to the coarser information set $\mathcal{H} \subseteq \mathcal{G}$. (Equivalently, think of the same econometrician at two points in time: at an early date she only knows \mathcal{H} , and later she learns the more detailed \mathcal{G} .)

The econometrician with information \mathcal{G} forms the forecast $\mathbb{E}[X | \mathcal{G}]$. The econometrician with less information cannot see \mathcal{G} directly, but she can try to forecast the forecast itself. Her prediction is $\mathbb{E}[\mathbb{E}[X | \mathcal{G}] | \mathcal{H}]$.

The tower property asserts that this indirect procedure gives exactly the same answer as if she had forecast X directly from \mathcal{H} :

$$\mathbb{E}[\mathbb{E}[X | \mathcal{G}] | \mathcal{H}] = \mathbb{E}[X | \mathcal{H}].$$

In other words, from the point of view of the coarser information set, it makes no difference whether you first condition on the finer information and then average, or condition directly.

Economically, this means that expectations are internally consistent across different levels of information. A forecast made at an early stage is the expected value of the more informed forecast that will be made later.

Theorem 4.8 (Taking out what is known). *If Z is \mathcal{G} -measurable and X is integrable, then*

$$\mathbb{E}[XZ | \mathcal{G}] = Z \mathbb{E}[X | \mathcal{G}].$$

Proof. Since Z is \mathcal{G} -measurable it is constant on each A_i : say $Z = z_i$ on A_i . Then

$$\mathbb{E}[XZ | A_i] = \frac{\int_{A_i} XZ dP}{P(A_i)} = z_i \cdot \frac{\int_{A_i} X dP}{P(A_i)} = z_i \mathbb{E}[X | A_i].$$

Hence the property holds on each atom and thus on all of Ω . □

Theorem 4.9 (Jensen's inequality). *If $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ is convex and X is integrable, then*

$$\varphi(\mathbb{E}[X | \mathcal{G}]) \leq \mathbb{E}[\varphi(X) | \mathcal{G}] \quad \text{pointwise on } \Omega.$$

4.7 Applications

A central use of conditional expectation in economics is forecasting. Suppose Y_{t+1} is an economic variable of interest (such as next period's output, inflation, or asset return) and \mathcal{F}_t is the information available at date t . The optimal forecast, in the mean-squared error sense, is

$$\mathbb{E}[Y_{t+1} \mid \mathcal{F}_t].$$

This expression captures the idea that agents use the information they have today to form beliefs about tomorrow.

In econometrics, this conditional expectation is the population regression function: given data Z , the best predictor of X is $\mathbb{E}[X \mid Z]$. In finance, it appears in the fundamental pricing relation: the current price of an asset is the conditional expectation of its discounted payoff, given the market's information set.

Thus conditional expectation provides a unified mathematical language for modeling expectations in macroeconomics, econometrics, and finance.

4.8 Exercises

Exercise 1. Let $\Omega = \{1, 2, \dots, 8\}$ with uniform probability $P(\{\omega\}) = 1/8$ for each ω . Define the random variable

$$X(1) = 0, X(2) = 1, X(3) = 2, X(4) = 3, X(5) = 4, X(6) = 5, X(7) = 6, X(8) = 7.$$

Consider the partitions

$$\mathcal{P} = \{\{1\}, \{2\}, \dots, \{8\}\}, \quad \mathcal{G} = \{\{1, 2\}, \{3, 4\}, \{5, 6\}, \{7, 8\}\},$$

and

$$\mathcal{H} = \{\{1, 2, 3, 4\}, \{5, 6, 7, 8\}\}.$$

Thus $\sigma(\mathcal{P})$ is the finest sigma-algebra, \mathcal{G} is a genuine coarsening of $\sigma(\mathcal{P})$, and \mathcal{H} is a coarsening of \mathcal{G} .

1. Compute $\mathbb{E}[X \mid \mathcal{G}]$ by evaluating the average of X on each atom of \mathcal{G} . Write the result in the form

$$\mathbb{E}[X \mid \mathcal{G}](\omega) = \sum_{i=1}^4 m_i \mathbf{1}_{A_i}(\omega),$$

where $A_1 = \{1, 2\}, A_2 = \{3, 4\}, A_3 = \{5, 6\}, A_4 = \{7, 8\}$ and give the numbers m_i .

2. Compute $\mathbb{E}[\mathbb{E}[X \mid \mathcal{G}] \mid \mathcal{H}]$ by averaging the values of $\mathbb{E}[X \mid \mathcal{G}]$ over the atoms of \mathcal{H} . Write the result as a two-step function on the atoms $\{1, 2, 3, 4\}$ and $\{5, 6, 7, 8\}$.

3. Compute $\mathbb{E}[X \mid \mathcal{H}]$ directly by averaging X on the atoms of \mathcal{H} .
4. Verify that

$$\mathbb{E}[\mathbb{E}[X \mid \mathcal{G}] \mid \mathcal{H}] = \mathbb{E}[X \mid \mathcal{H}],$$

thereby confirming the tower property for this genuine coarsening.

Exercise 2. Let $\Omega = \{a, b, c\}$ with probabilities $P(\{a\}) = 1/2$, $P(\{b\}) = 1/3$, $P(\{c\}) = 1/6$. Consider the partition $\mathcal{P} = \{A_1, A_2\}$ with $A_1 = \{a, b\}$, $A_2 = \{c\}$ and let $\mathcal{G} = \sigma(\mathcal{P})$. Define

$$X(a) = 2, X(b) = 1, X(c) = 0, \quad Z(a) = 3, Z(b) = 3, Z(c) = 5.$$

Note that Z is \mathcal{G} -measurable (constant on each atom). Compute both sides of the identity

$$\mathbb{E}[XZ \mid \mathcal{G}] \stackrel{?}{=} Z \mathbb{E}[X \mid \mathcal{G}]$$

by evaluating each atom of \mathcal{P} . Verify the equality numerically.

Exercise 3.

- (a) Discrete case. Let (X, Y) have the joint mass function

	$Y = 0$	$Y = 1$
$X = 0$	0.2	0.1
$X = 1$	0.3	0.4

(rows sum to 1). For $y = 0, 1$ compute the conditional pmf $P(X = x \mid Y = y)$ and then compute $\mathbb{E}[X \mid Y = y]$. Write the random variable $\mathbb{E}[X \mid Y]$ as a function of Y .

- (b) Continuous case. Let (X, Y) have joint density

$$f_{X,Y}(x, y) = 2, \quad (x, y) \in \{0 \leq x \leq y \leq 1\},$$

and $f_{X,Y}(x, y) = 0$ elsewhere. Compute the marginal density $f_Y(y)$, the conditional density $f_{X|Y}(x \mid y)$ for $0 \leq x \leq y$, and then compute

$$\mathbb{E}[X \mid Y = y] = \int_0^y x f_{X|Y}(x \mid y) dx.$$

Express $\mathbb{E}[X \mid Y]$ as a function of Y .

5 Martingales

5.1 Filtered Probability Spaces and Adapted Processes

We begin by recalling the basic setting in which martingale theory is formulated.

Definition 5.1 (Filtration). A *filtration* on a probability space (Ω, \mathcal{F}, P) is a sequence of sub-sigma-algebras

$$\mathcal{G}_0 \subseteq \mathcal{G}_1 \subseteq \cdots \subseteq \mathcal{F}.$$

Each \mathcal{G}_n represents the information available up to time n .

Definition 5.2 (Adapted process). A stochastic process $(X_n)_{n \geq 0}$ is said to be *adapted* to the filtration (\mathcal{G}_n) if X_n is \mathcal{G}_n -measurable for every n .

5.2 Martingales, Submartingales, and Supermartingales

Definition 5.3 (Martingale, submartingale, supermartingale). Let $(X_n)_{n \geq 0}$ be an adapted process with $\mathbb{E}[|X_n|] < \infty$ for all n (i.e. integrable).

- X is a *martingale* if

$$\mathbb{E}[X_{n+1} \mid \mathcal{G}_n] = X_n \quad \text{a.s. for all } n.$$

- X is a *submartingale* if

$$\mathbb{E}[X_{n+1} \mid \mathcal{G}_n] \geq X_n.$$

- X is a *supermartingale* if

$$\mathbb{E}[X_{n+1} \mid \mathcal{G}_n] \leq X_n.$$

Gambling story (Motivation)

Think of X_n as the total wealth of a gambler, who bets one Rupee per game, after n games. The increment

$$X_{n+1} - X_n$$

represents the profit or loss incurred in the $(n+1)$ -st game.

The n th game is called *fair* if

$$\mathbb{E}[X_{n+1} - X_n \mid \mathcal{G}_n] = 0,$$

unfair against the gambler if

$$\mathbb{E}[X_{n+1} - X_n \mid \mathcal{G}_n] \leq 0,$$

and *favorable* if

$$\mathbb{E}[X_{n+1} - X_n \mid \mathcal{G}_n] \geq 0.$$

5.3 Basic Examples

Proposition 5.4 (Three canonical martingales). *All statements are with respect to the natural filtration $\mathcal{G}_n = \sigma(Y_1, \dots, Y_n)$ or $\sigma(Z_1, \dots, Z_n)$ (as in Definition 1.4).*

- (a) **Sum of independent zero-mean variables.** *Let $(Y_k)_{k \geq 1}$ be independent with $\mathbb{E}[|Y_k|] < \infty$ and $\mathbb{E}[Y_k] = 0$. Define*

$$M_n := \sum_{k=1}^n Y_k, \quad n \geq 1, \quad M_0 := 0.$$

Then (M_n) is a martingale with respect to $\mathcal{G}_n = \sigma(Y_1, \dots, Y_n)$.

- (b) **Product of independent unit-mean variables (no nonnegativity needed).** *Let $(Z_k)_{k \geq 1}$ be independent with $\mathbb{E}[|Z_k|] < \infty$ and $\mathbb{E}[Z_k] = 1$ for all k . Define*

$$M_n := \prod_{k=1}^n Z_k, \quad n \geq 1, \quad M_0 := 1,$$

and let $\mathcal{G}_n = \sigma(Z_1, \dots, Z_n)$. Then (M_n) is a martingale.

- (c) **Conditional expectation process.** *Let X be integrable and $(\mathcal{G}_n)_{n \geq 0}$ a filtration. Define*

$$M_n := \mathbb{E}[X \mid \mathcal{G}_n], \quad n \geq 0.$$

Then (M_n) is a martingale with respect to (\mathcal{G}_n) .

Proof. (a) *Adaptedness/integrability:* M_n is \mathcal{G}_n -measurable as a sum of \mathcal{G}_n -measurable random variables; $\mathbb{E}|M_n| \leq \sum_{k=1}^n \mathbb{E}|Y_k| < \infty$. *Martingale step:*

$$\mathbb{E}[M_{n+1} \mid \mathcal{G}_n] = \mathbb{E}[M_n + Y_{n+1} \mid \mathcal{G}_n] = M_n + \mathbb{E}[Y_{n+1} \mid \mathcal{G}_n].$$

By independence, Y_{n+1} is independent of \mathcal{G}_n , so $\mathbb{E}[Y_{n+1} \mid \mathcal{G}_n] = \mathbb{E}[Y_{n+1}] = 0$. Hence $\mathbb{E}[M_{n+1} \mid \mathcal{G}_n] = M_n$ a.s.

- (b) *Adaptedness/integrability:* M_n is \mathcal{G}_n -measurable. By independence, $\mathbb{E}[|M_n|] \leq \prod_{k=1}^n \mathbb{E}[|Z_k|] < \infty$, so M_n is integrable; likewise M_{n+1} . Moreover,

$$\mathbb{E}[M_{n+1} \mid \mathcal{G}_n] = \mathbb{E}[M_n Z_{n+1} \mid \mathcal{G}_n] = M_n \mathbb{E}[Z_{n+1}] = M_n.$$

Hence (M_n) is a martingale.

- (c) *Adaptedness/integrability:* By definition, M_n is \mathcal{G}_n -measurable and $\mathbb{E}|M_n| \leq \mathbb{E}|X| < \infty$. *Martingale step (tower property):*

$$\mathbb{E}[M_{n+1} \mid \mathcal{G}_n] = \mathbb{E}[\mathbb{E}[X \mid \mathcal{G}_{n+1}] \mid \mathcal{G}_n] = \mathbb{E}[X \mid \mathcal{G}_n] = M_n \quad \text{a.s.}$$

□

5.4 Previsible Processes and Martingale Transforms

We now introduce the notion of predictable processes and use them to define linear transformations of martingales. These objects provide a rigorous formalization of strategies in stochastic models.

Definition 5.5 (Predictable process). Let $(\mathcal{G}_n)_{n \geq 0}$ be a filtration. A stochastic process $(C_n)_{n \geq 1}$ is said to be *predictable* (with respect to (\mathcal{G}_n)) if

$$C_n \text{ is } \mathcal{G}_{n-1}\text{-measurable for each } n \geq 1.$$

Equivalently, the choice of C_n may depend on the information available up to and including time $n - 1$, but not on the outcome of period n itself.

Definition 5.6 (Martingale transform with initial capital). Let $X = (X_n)_{n \geq 0}$ be an integrable process and $C = (C_n)_{n \geq 1}$ a predictable process. For an initial capital Y_0 with $Y_0 \in L^1$ and Y_0 \mathcal{G}_0 -measurable, define

$$(C \cdot X)_n := \sum_{k=1}^n C_k (X_k - X_{k-1}), \quad n \geq 1, \quad (C \cdot X)_0 := 0,$$

and set the transformed wealth process

$$Y_n := Y_0 + (C \cdot X)_n, \quad n \geq 0.$$

Theorem 5.7 (Stability under bounded predictable transforms). *Let $X = (X_n)_{n \geq 0}$ be a supermartingale (resp. martingale) with respect to (\mathcal{G}_n) . Let $C = (C_n)_{n \geq 1}$ be predictable and assume there exists $K < \infty$ such that*

$$|C_n| \leq K \quad \text{almost surely for all } n \geq 1.$$

Let Y_0 be \mathcal{G}_0 -measurable and integrable, and define

$$Y_n := Y_0 + \sum_{k=1}^n C_k (X_k - X_{k-1}), \quad n \geq 0.$$

Then (Y_n) is a supermartingale (resp. martingale) with respect to (\mathcal{G}_n) .

Proof. Each increment $C_k(X_k - X_{k-1})$ is \mathcal{G}_k -measurable, hence Y_n is \mathcal{G}_n -measurable. We check integrability:

$$\begin{aligned} \mathbb{E}[|Y_n|] &\leq \mathbb{E}[|Y_0|] + \sum_{k=1}^n \mathbb{E}[|C_k| |X_k - X_{k-1}|] \\ &\leq \mathbb{E}[|Y_0|] + K \sum_{k=1}^n \mathbb{E}[|X_k| + |X_{k-1}|] < \infty, \end{aligned}$$

since each X_m is integrable. Thus $Y_n \in L^1$.

Now compute the conditional expectation. For $n \geq 0$,

$$\mathbb{E}[Y_{n+1} \mid \mathcal{G}_n] = Y_n + \mathbb{E}[C_{n+1}(X_{n+1} - X_n) \mid \mathcal{G}_n].$$

By previsibility, C_{n+1} is \mathcal{G}_n -measurable, so

$$\mathbb{E}[Y_{n+1} \mid \mathcal{G}_n] = Y_n + C_{n+1} \mathbb{E}[X_{n+1} - X_n \mid \mathcal{G}_n].$$

If X is a supermartingale, the conditional increment is ≤ 0 , hence $\mathbb{E}[Y_{n+1} \mid \mathcal{G}_n] \leq Y_n$. If X is a martingale, the inequality is an equality. \square

Gambling story (Consequence)

Interpret (X_n) as the wealth of a unit-stake gambler after n fair games, so (X_n) is a martingale. A previsible process (C_n) represents the gambler's strategy for choosing stakes: before each game, the amount C_n to be bet is decided on the basis of past information \mathcal{G}_{n-1} .

If the gambler's stakes are bounded (which in practice means the gambler has only finite capital to risk) then the transformed wealth

$$Y_n = Y_0 + \sum_{k=1}^n C_k (X_k - X_{k-1})$$

remains a martingale.

As long as the gambler cannot stake arbitrarily large amounts, no betting strategy can convert fair games into a source of systematic profit. The “fairness” of the casino is preserved under all feasible strategies.

5.5 Stopping Times and Stopped Processes

Throughout this section, for two integers m, n we write $m \wedge n := \min\{m, n\}$.

Definition 5.8 (Stopping time). Let $(\mathcal{G}_n)_{n \geq 0}$ be a filtration on $(\Omega, \mathcal{F}, \mathbb{P})$. A random variable $T : \Omega \rightarrow \mathbb{N} \cup \{\infty\}$ is called a *stopping time* with respect to (\mathcal{G}_n) if

$$\{T \leq n\} \in \mathcal{G}_n \quad \text{for every } n \geq 0.$$

In other words, by time n it is observable whether or not the stopping time has occurred.

Definition 5.9 (Stopped process). Let $X = (X_n)_{n \geq 0}$ be a stochastic process and let T be a stopping time. The *stopped process* $Z = (Z_n)_{n \geq 0}$ is defined by

$$Z_n := X_{n \wedge T}, \quad n \geq 0,$$

so that Z coincides with X up to time T and remains constant thereafter.

Theorem 5.10 (Stopped (super)martingales under bounded horizons). *Let $(\mathcal{G}_n)_{n \geq 0}$ be a filtration and let $X = (X_n)_{n \geq 0}$ be a supermartingale (resp. martingale) with respect to (\mathcal{G}_n) . Let T be a bounded stopping time, i.e. there exists $N < \infty$ with $T \leq N$ almost surely, and define the stopped process $Z = (Z_n)_{n \geq 0}$ by*

$$Z_n := X_{n \wedge T}, \quad n \geq 0.$$

Then the following holds:

- (i) *For each n , the random variable Z_n is integrable, i.e. $\mathbb{E}[|Z_n|] < \infty$.*
- (ii) *The process Z is a supermartingale (resp. martingale) with respect to (\mathcal{G}_n) .*

Proof. Since $T \leq N$ almost surely, for each n we have

$$Z_n = X_{n \wedge T} \in \{X_0, X_1, \dots, X_{\min\{n, N\}}\}.$$

This means Z_n takes values among finitely many random variables, each of which is integrable because X is a (super)martingale. Hence $\mathbb{E}[|Z_n|] < \infty$, proving (i). Adaptedness of Z follows immediately from the definition.

For (ii), observe that

$$\mathbb{E}[Z_{n+1} \mid \mathcal{G}_n] = \mathbb{E}[X_{(n+1) \wedge T} \mid \mathcal{G}_n].$$

If $T \leq n$, then $(n+1) \wedge T = T = n \wedge T$ and so $X_{(n+1) \wedge T} = Z_n$. If $T > n$, then $(n+1) \wedge T = n+1$ and

$$\mathbb{E}[X_{(n+1) \wedge T} \mid \mathcal{G}_n] = \mathbb{E}[X_{n+1} \mid \mathcal{G}_n].$$

Combining the two cases gives

$$\mathbb{E}[Z_{n+1} \mid \mathcal{G}_n] = Z_n \mathbf{1}_{\{T \leq n\}} + \mathbb{E}[X_{n+1} \mid \mathcal{G}_n] \mathbf{1}_{\{T > n\}}.$$

If X is a supermartingale then $\mathbb{E}[X_{n+1} \mid \mathcal{G}_n] \leq X_n$, and since $Z_n = X_n$ on $\{T > n\}$ it follows that $\mathbb{E}[Z_{n+1} \mid \mathcal{G}_n] \leq Z_n$. If X is a martingale, equality holds. Thus Z inherits the (super)martingale property, establishing (ii). \square

Gambling story (Consequence)

Interpret X_n as the wealth of a gambler after n rounds, in a (super)fair game. A stopping time T models a rule such as “stop once my fortune reaches a target or I fall below a limit,” with the rule depending only on past outcomes. The theorem says that if such a rule has a bounded horizon, then the stopped wealth process $Z_n = X_{n \wedge T}$ preserves the (super)martingale property. In plain terms: even if the gambler chooses to stop according to a legitimate strategy, no advantage can be gained on average—the game remains fair (or unfavorable) up to that stopping time.

5.5.1 Applications of Optional Sampling

One of the most important uses of the optional sampling theorem is in the analysis of hitting times for random walks. Let $(S_n)_{n \geq 0}$ be a simple symmetric random walk, defined by $S_0 = 0$ and $S_n = \sum_{k=1}^n W_k$, where the increments (W_k) are i.i.d. with $\mathbb{P}(W_k = 1) = \mathbb{P}(W_k = -1) = 1/2$. The process (S_n) is a martingale with respect to its natural filtration. Suppose we fix two integers $a < 0 < b$ and define the stopping time

$$T := \inf\{n \geq 0 : S_n = a \text{ or } S_n = b\}.$$

That is, T is the first time the walk leaves the open interval (a, b) . Since T is bounded above by the finite number $|a| + b$, it is a bounded stopping time. The optional sampling theorem therefore guarantees that

$$\mathbb{E}[S_T] = \mathbb{E}[S_0] = 0.$$

Because S_T takes only the values a and b , we obtain

$$0 = \mathbb{E}[S_T] = a \cdot \mathbb{P}(S_T = a) + b \cdot \mathbb{P}(S_T = b).$$

Solving for the probabilities gives

$$\mathbb{P}(S_T = b) = \frac{|a|}{|a| + b}, \quad \mathbb{P}(S_T = a) = \frac{b}{|a| + b}.$$

This calculation is the classical gambler's ruin formula: if a gambler begins with zero capital, plays fair games of ± 1 , and stops upon reaching either a units of debt or b units of profit, then the probability of ultimate success is proportional to the distance to ruin.

A second application concerns the expected duration of such fair games. Consider the same random walk and stopping time T , but now examine the process $(S_n^2 - n)_{n \geq 0}$. Because the increments W_k are independent, mean zero, and variance one, the process $(S_n^2 - n)$ is itself a martingale. To see this, recall that (W_k) are independent with $\mathbb{E}[W_k] = 0$ and $\mathbb{E}[W_k^2] = 1$. Write $S_{n+1} = S_n + W_{n+1}$. Then

$$S_{n+1}^2 - (n+1) = S_n^2 + 2S_n W_{n+1} + W_{n+1}^2 - (n+1).$$

Taking conditional expectation with respect to \mathcal{G}_n and using the fact that W_{n+1} is independent of \mathcal{G}_n ,

$$\mathbb{E}[S_{n+1}^2 - (n+1) \mid \mathcal{G}_n] = S_n^2 + 2S_n \mathbb{E}[W_{n+1}] + \mathbb{E}[W_{n+1}^2] - (n+1).$$

Since $\mathbb{E}[W_{n+1}] = 0$ and $\mathbb{E}[W_{n+1}^2] = 1$, this reduces to

$$\mathbb{E}[S_{n+1}^2 - (n+1) \mid \mathcal{G}_n] = S_n^2 - n = M_n.$$

Hence $(M_n)_{n \geq 0}$ is a martingale. Applying the optional sampling theorem at T yields

$$\mathbb{E}[S_T^2 - T] = \mathbb{E}[S_0^2 - 0] = 0.$$

Since S_T equals a or b , we have $\mathbb{E}[S_T^2] = a^2\mathbb{P}(S_T = a) + b^2\mathbb{P}(S_T = b)$. Substituting the ruin probabilities computed above leads to an explicit expression for $\mathbb{E}[T]$. A short calculation shows

$$\mathbb{E}[T] = |a|b.$$

These examples illustrate the general principle: the combination of martingale structure and stopping times allows one to compute both hitting probabilities and expected hitting times for random walks. The method requires no delicate combinatorics and avoids recursive difference equations, relying only on the fundamental identity $\mathbb{E}[X_T] = \mathbb{E}[X_0]$ when X is a martingale and T is a bounded stopping time. This simplicity is precisely what makes martingale theory indispensable in probability and in its economic applications.

5.5.2 Important Non-Example: the doubling (“martingale”) strategy

Let $(W_n)_{n \geq 1}$ be an i.i.d. sequence with $\mathbb{P}(W_n = 1) = \mathbb{P}(W_n = -1) = 1/2$. The simple random walk $S_n = \sum_{k=1}^n W_k$ is then a martingale with respect to its natural filtration.

Consider now the following strategy. Define the stopping time

$$T := \min\{n \geq 1 : W_n = 1\},$$

and set the previsible stakes

$$C_n := 2^{n-1} \mathbf{1}_{\{W_1 = \dots = W_{n-1} = -1\}}, \quad n \geq 1.$$

The gambler doubles her stake after each loss and ceases to bet once the first win has occurred. The corresponding transformed wealth is

$$Y_n := \sum_{k=1}^{T \wedge n} C_k W_k, \quad n \geq 0.$$

On the event $\{T \leq n\}$ the first win occurs by time n . In this case the gambler loses $1, 2, 4, \dots, 2^{T-2}$ on the successive losses and finally gains 2^{T-1} on the win, so that $Y_n = 1$. On the complementary event $\{T > n\}$ all of the first n outcomes are losses, and the gambler’s wealth is

$$Y_n = -(2^n - 1).$$

It follows that for every finite n ,

$$\mathbb{E}[Y_n] = 1 \cdot (1 - 2^{-n}) + (-(2^n - 1)) \cdot 2^{-n} = 0.$$

Thus the doubling strategy preserves mean zero at each finite stage. However, since $T < \infty$ almost surely, the sequence (Y_n) converges almost surely to $Y_T = 1$. Hence

$$\mathbb{E}[Y_T] = 1 \quad \text{while} \quad \lim_{n \rightarrow \infty} \mathbb{E}[Y_n] = 0.$$

This construction demonstrates that one cannot, in general, interchange limit and expectation. At finite horizons the strategy yields no expected gain, yet at the random horizon T the expectation has jumped to one. The apparent paradox arises because the doubling strategy requires unbounded stakes: each new loss forces the gambler to risk twice as much capital. Mathematically, this breaks the hypotheses under which optional sampling guarantees preservation of the martingale property.

5.6 Monkeys typing ABRACADABRA

We imagine a monkey typing letters independently and uniformly at random from the 26-letter alphabet. We ask for the expected number of keystrokes until the sequence of letters

ABRACADABRA

appears for the first time as a contiguous block. Let T denote this random time. Our goal is to compute $\mathbb{E}[T]$.

This problem has evoked the curiosity of a unit stake gambler and a gambling house. The gambler believes deeply that the string will appear and the house does not. The gambler starts a new series of bets at every key stroke (until the string appears). At every key stroke he bets A appears. If he is right, he wins 26 rupees and continues his series to bet that the next letter is B but this time he bets his 26 rupees. If he wins, then he receives 26^2 rupees and so on. He continues his series of bets until he loses for the first time. Note that if the game ends at time T , then he has bet a total of T rupees across T series² of bets. Now a thorough analysis using optional sampling theorem yields the answer. Read on.

Mathematically precise setup: Let $(X_n)_{n \geq 1}$ be i.i.d. uniform on the 26-letter alphabet, and let

$$T := \inf\{n \geq 11 : (X_{n-10}, \dots, X_n) = \text{ABRACADABRA}\}.$$

For each start time $s \geq 0$ consider the *run* that begins at s : it stakes 1 at time s that $X_{s+1} = \text{A}$; if correct, the stake is multiplied by 26 and is then staked on $X_{s+2} = \text{B}$; if correct again, the stake is multiplied by 26 and staked on $X_{s+3} = \text{R}$; and so on, always wagering the *entire current fortune* on the next required letter of **ABRACADABRA**. A wrong letter ends the run at 0; k consecutive correct letters produce capital 26^k . Thus each run is a sequence of *sequential* fair bets with multiplier 26 on success. A fresh run starts at every time s , so runs overlap in calendar time, but *within* each run the bets are strictly sequential.

²Not the record company that messed with pewdiepie lol!

From the house's perspective, define its fortune

$$H_n = H_0 + (\text{total stakes received up to time } n) - (\text{total payouts made up to time } n).$$

Each stake is 1, each per-letter wager is fair, and runs are independent of the house's accounting, so (H_n) is a martingale with respect to the natural filtration. To meet the worst case that the word appears immediately (in the first 11 letters), the house must be capitalized by

$$H_0 = P - 11, \quad P := 26^{11} + 26^4 + 26,$$

because on that event it will have received 11 units of stakes and must pay the unique surviving runs: the length-11 run (full word), the length-4 run (suffix/prefix "ABRA"), and the length-1 run (suffix/prefix "A").

At the stopping time T , the house has received exactly T units of stakes and must pay precisely P , since the only runs alive at T are those whose starts align with the borders of ABRACADABRA: the run from $T - 10$ has capital 26^{11} , the run from $T - 3$ has capital 26^4 , and the run from T has capital 26. Hence

$$H_T = H_0 + T - P = T - 11.$$

By the optional stopping theorem,

$$\mathbb{E}[H_T] = \mathbb{E}[H_0],$$

so $\mathbb{E}[T - 11] = P - 11$ and therefore

$$\boxed{\mathbb{E}[T] = 26^{11} + 26^4 + 26}.$$

5.7 Markov Chains and Martingales

Let $X = (X_n)_{n \geq 0}$ be a time-homogeneous Markov chain with state space S and transition matrix $P = (P(i, j))_{i, j \in S}$, defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{G}_n), \mathbb{P})$ with natural filtration $\mathcal{G}_n = \sigma(X_0, \dots, X_n)$.

Proposition 5.11. *If $f : S \rightarrow \mathbb{R}$ satisfies*

$$f(i) = \sum_{j \in S} P(i, j) f(j), \quad i \in S,$$

then $M_n := f(X_n)$ is a martingale with respect to (\mathcal{G}_n) .

Proof. For each $n \geq 0$,

$$\mathbb{E}[M_{n+1} \mid \mathcal{G}_n] = \mathbb{E}[f(X_{n+1}) \mid X_n] = \sum_{j \in S} P(X_n, j) f(j).$$

If $X_n = i$, this equals $\sum_j P(i, j) f(j) = f(i) = M_n$. Hence M is a martingale. \square

A function f satisfying the relation above is called *harmonic* with respect to P . Thus every harmonic function gives rise to a martingale along the paths of the chain.

5.7.1 Explicit Harmonic Examples and Optional stopping calculations

Example 5.12 (Symmetric random walk: exit probabilities on $\{0, 1, \dots, N\}$). Consider the simple symmetric random walk $X = (X_n)_{n \geq 0}$ on the finite state space $\{0, 1, \dots, N\}$ with absorbing endpoints 0 and N . That is, for $1 \leq i \leq N - 1$,

$$\mathbb{P}(X_{n+1} = i + 1 \mid X_n = i) = \mathbb{P}(X_{n+1} = i - 1 \mid X_n = i) = \frac{1}{2},$$

while for $i \in \{0, N\}$ one has $X_{n+1} = X_n = i$ almost surely.

Define the function

$$f(i) := \frac{i}{N}, \quad 0 \leq i \leq N.$$

Then for every interior state $1 \leq i \leq N - 1$,

$$\mathbb{E}[f(X_{n+1}) \mid X_n = i] = \frac{1}{2}f(i + 1) + \frac{1}{2}f(i - 1) = \frac{i}{N} = f(i).$$

Thus the process $M_n := f(X_n)$ is a martingale with respect to the natural filtration of the chain.

Let

$$T := \inf\{n \geq 0 : X_n \in \{0, N\}\}$$

denote the first hitting time of the boundary. Since $0 \leq T \leq N$ almost surely, the stopping time T is bounded. The bounded optional sampling theorem therefore applies, and gives

$$\mathbb{E}[f(X_T) \mid X_0 = x] = f(x), \quad 0 \leq x \leq N.$$

At the stopping time T the chain is absorbed, so $X_T \in \{0, N\}$. Consequently,

$$f(X_T) = \frac{X_T}{N} = \begin{cases} 0, & X_T = 0, \\ 1, & X_T = N, \end{cases} = \mathbf{1}_{\{X_T = N\}}.$$

Hence

$$\mathbb{P}(X_T = N \mid X_0 = x) = \mathbb{E}[\mathbf{1}_{\{X_T = N\}} \mid X_0 = x] = \mathbb{E}[f(X_T) \mid X_0 = x] = f(x) = \frac{x}{N}.$$

It follows that

$$\mathbb{P}(X_T = N \mid X_0 = x) = \frac{x}{N}, \quad \mathbb{P}(X_T = 0 \mid X_0 = x) = 1 - \frac{x}{N}.$$

Example 5.13 (Biased random walk: gambler's ruin probabilities). Consider the nearest-neighbour random walk $X = (X_n)_{n \geq 0}$ on $\{0, 1, \dots, N\}$ with absorbing states 0 and N . For each interior state

$$1 \leq i \leq N-1,$$

$$\mathbb{P}(X_{n+1} = i+1 \mid X_n = i) = p, \quad \mathbb{P}(X_{n+1} = i-1 \mid X_n = i) = q := 1-p,$$

while $X_{n+1} = X_n$ when $X_n \in \{0, N\}$. Assume throughout that $p \neq q$.

Define

$$f(i) := r^i, \quad 0 \leq i \leq N, \quad r := \frac{q}{p}.$$

Verification of the martingale property. For each interior state $1 \leq i \leq N-1$,

$$\mathbb{E}[f(X_{n+1}) \mid X_n = i] = p r^{i+1} + q r^{i-1} = r^i(pr + q/r) = r^i = f(i).$$

At the boundaries, if $X_n = 0$ then $X_{n+1} = 0$ almost surely, so $f(X_{n+1}) = f(0) = 1 = f(X_n)$. Similarly, if $X_n = N$ then $X_{n+1} = N$ almost surely, so $f(X_{n+1}) = f(N) = r^N = f(X_n)$. Thus

$$\mathbb{E}[f(X_{n+1}) \mid \mathcal{G}_n] = f(X_n)$$

for all n , and therefore $M_n := f(X_n)$ is a martingale.

Application of optional sampling. Let

$$T := \inf\{n \geq 0 : X_n \in \{0, N\}\}$$

be the first hitting time of the boundary. Since $0 \leq T \leq N$ almost surely, T is a bounded stopping time. The optional sampling theorem then yields

$$\mathbb{E}[f(X_T) \mid X_0 = x] = f(x) = r^x, \quad 0 \leq x \leq N.$$

At the stopping time T the chain is absorbed, so $X_T \in \{0, N\}$. Consequently

$$f(X_T) = \begin{cases} 1, & X_T = 0, \\ r^N, & X_T = N. \end{cases}$$

Therefore

$$r^x = \mathbb{E}[f(X_T) \mid X_0 = x] = \mathbb{P}(X_T = 0 \mid X_0 = x) \cdot 1 + \mathbb{P}(X_T = N \mid X_0 = x) \cdot r^N.$$

Solving for the probabilities gives

$$\mathbb{P}(X_T = N \mid X_0 = x) = \frac{1 - r^x}{1 - r^N}, \quad \mathbb{P}(X_T = 0 \mid X_0 = x) = \frac{r^x - r^N}{1 - r^N}.$$

In the limit $p \rightarrow q$ (so $r \rightarrow 1$), these probabilities reduce to those of Example 5.12, namely $\mathbb{P}(X_T = N \mid X_0 = x) = x/N$.

5.8 Exercises

Exercise 5.14. Let $(X_i)_{i \geq 1}$ be independent random variables with finite expectations and variances,

$$\mathbb{E}[X_i] = m_i, \quad \text{Var}(X_i) = \sigma_i^2,$$

and set $S_n = \sum_{i=1}^n X_i$, $\mathcal{F}_n = \sigma(X_1, \dots, X_n)$.

1. Find sequences (b_n) and (c_n) of real numbers such that

$$S_n^2 + b_n S_n + c_n$$

is an (\mathcal{F}_n) -martingale.

2. Fix $\lambda \in \mathbb{R}$ and assume $\mathbb{E}[e^{\lambda X_i}] < \infty$ for all i . Define $G_i(\lambda) := \mathbb{E}[e^{\lambda X_i}]$. Find a sequence $(a_n^\lambda)_{n \geq 0}$ such that

$$\exp(\lambda S_n - a_n^\lambda)$$

is an (\mathcal{F}_n) -martingale.

Exercise 5.15. Let $X = (X_n)_{n \geq 0}$ be the simple symmetric random walk on \mathbb{Z} with natural filtration (\mathcal{F}_n) .

1. Show that the process

$$Y_n := X_n^3 - 3nX_n, \quad n \geq 0,$$

is an (\mathcal{F}_n) -martingale.

2. Fix $N \in \mathbb{N}$ and let

$$\tau := \inf\{n \geq 0 : X_n \in \{0, N\}\}$$

be the first hitting time of the boundary $\{0, N\}$. For $0 \leq k \leq N$, compute

$$\mathbb{E}_k[\tau \mid X_\tau = N],$$

where \mathbb{E}_k denotes expectation given $X_0 = k$.

Exercise 5.16. Let $(X_n)_{n \geq 0}$ be a martingale with respect to (\mathcal{F}_n) .

1. Show that its increments are pairwise orthogonal, i.e.

$$\mathbb{E}[(X_{n+1} - X_n)(X_{m+1} - X_m)] = 0 \quad \text{for all } n \neq m.$$

2. Suppose in addition that the increments $X_{n+1} - X_n$ are $\{0, 1\}$ -valued or Gaussian. Show that (X_n) is a random walk.

Exercise 5.17 (Waiting time for coin patterns). We toss a fair coin repeatedly.

1. Let T be the first time the sequence HTHT appears consecutively. Show that T is almost surely finite.
2. Consider the process $X_n = 2^{L_n}$, where L_n is the length of the longest suffix of the first n tosses that is also a prefix of HTHT. Show that (X_n) is a martingale.
3. Use the martingale (X_n) and the optional stopping theorem at T to compute $\mathbb{E}[T]$.
4. Among all 16 patterns of length 4, give an example of one with the maximal expected waiting time. (Hint: compare patterns with and without overlaps.)

6 Martingales in Financial Markets

Financial markets are built on uncertainty. Prices move with news, information, and chance. The question we face is: how can such random movements be understood mathematically, and how do we decide what a fair price for a risky asset should be?

Two results form the backbone of modern asset pricing:

- The first is the *fundamental pricing equation*

$$p_t = \mathbb{E}_t[m_{t+1}x_{t+1}],$$

which says that today's price of an asset equals the conditional expectation of its payoff tomorrow, weighted by a random variable m_{t+1} called the *stochastic discount factor*. This single formula captures both the time value of money and the adjustment for risk.

- The second is the *Fundamental Theorem of Asset Pricing (FTAP)*. It tells us that the simple requirement of “no arbitrage” is exactly the same as the existence of a probability measure under which discounted asset prices are martingales. In other words, if there are no free lunches in the market, then prices evolve as fair games once we account for the numéraire.

The first equation comes from economics: it is the condition that arises when agents with preferences over consumption choose optimally between today and tomorrow. The second comes from probability and convex geometry: it shows that the martingale property of discounted prices is forced by the absence of arbitrage.

The aim of this chapter is to connect these two points of view. We will move in small steps:

1. Start with risk preferences and utility, using Jensen's inequality to explain risk aversion, neutrality, and seeking.
2. Derive the pricing equation $p_t = \mathbb{E}_t[m_{t+1}x_{t+1}]$ from a simple two-period optimisation problem.
3. Work through examples such as bonds and stocks with dividends.
4. Introduce convex sets and a separation result, which provide the mathematical language needed for arbitrage arguments.
5. State and prove the Fundamental Theorem of Asset Pricing in discrete time, showing how martingales enter naturally.

In this way we pass from economic intuition about preferences to the precise mathematical structure that underlies modern finance.

6.1 Risk Preferences and Utility

We begin by recalling how preferences over random payoffs can be described. The basic objects of choice are *lotteries*, that is, random variables representing uncertain consumption or wealth.

6.1.1 Definitions

Definition 6.1 (Lottery). A *lottery* is a random variable X on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, taking values in \mathbb{R}_+ and representing a nonnegative payoff.

Definition 6.2 (Expected Utility Representation). Let \succeq denote the agent's preference relation on lotteries, where $X \succeq Y$ means that the lottery X is *weakly preferred* to the lottery Y (the agent likes X at least as much as Y).

An agent is said to have an *expected utility representation* if there exists a function $u : \mathbb{R}_+ \rightarrow \mathbb{R}$, called the *utility function*, such that for any two lotteries X, Y one has

$$X \succeq Y \iff \mathbb{E}[u(X)] \geq \mathbb{E}[u(Y)].$$

Throughout this chapter, we will assume an expected utility representation exists for the agents we are interested in³.

Definition 6.3 (Risk Preferences). Let X be a lottery with mean $\mu = \mathbb{E}[X]$.

1. The agent is *risk averse* if

$$u(\mu) \geq \mathbb{E}[u(X)].$$

In words: the utility of receiving the mean payoff for sure is at least as large as the expected utility of the risky payoff.

2. The agent is *risk neutral* if

$$u(\mu) = \mathbb{E}[u(X)],$$

i.e. the agent cares only about the mean payoff, not about the distribution around it.

3. The agent is *risk seeking* if

$$u(\mu) \leq \mathbb{E}[u(X)].$$

That is, the agent prefers risk to certainty with the same expected payoff.

For whatever follows, we need Jensen's inequality. We state it here instead of the appendix because this inequality is of independent interest.

Theorem 6.4 (Jensen's Inequality). *Let $u : \mathbb{R} \rightarrow \mathbb{R}$ be convex, and let X be an integrable random variable. Then*

$$u(\mathbb{E}[X]) \leq \mathbb{E}[u(X)].$$

If u is concave, the inequality is reversed:

$$u(\mathbb{E}[X]) \geq \mathbb{E}[u(X)].$$

³The fact that such a representation exists is the content of the *von Neumann–Morgenstern utility theorem*. It states that if an agent's preferences over lotteries satisfy certain axioms (completeness, transitivity, continuity, and independence), then there exists a utility function u unique up to positive affine transformations such that the preferences can be written in the expected utility form above.

So now we can connect the risk taking behaviour to the curvature of the utility function.

Proposition 6.5. *Let X be a lottery with finite expectation, and suppose an agent's preferences are represented by an expected utility function $u : \mathbb{R}_+ \rightarrow \mathbb{R}$. Then:*

1. *If u is concave, the agent is risk averse, i.e.*

$$u(\mathbb{E}[X]) \geq \mathbb{E}[u(X)].$$

2. *If u is affine (linear up to a constant), the agent is risk neutral, i.e.*

$$u(\mathbb{E}[X]) = \mathbb{E}[u(X)].$$

3. *If u is convex, the agent is risk seeking, i.e.*

$$u(\mathbb{E}[X]) \leq \mathbb{E}[u(X)].$$

Proof. This follows directly from Jensen's inequality. If u is concave, $u(\mathbb{E}[X]) \geq \mathbb{E}[u(X)]$, which is precisely the definition of risk aversion. If u is affine, Jensen's inequality holds with equality, showing risk neutrality. If u is convex, the inequality reverses, showing risk seeking. \square

Example 6.6 (Lottery vs. Sure Payoff). Consider a lottery that pays

$$X = \begin{cases} 100 & \text{with probability } \frac{1}{2}, \\ 0 & \text{with probability } \frac{1}{2}. \end{cases}$$

The expected value is $\mathbb{E}[X] = 50$. Compare this lottery to the sure payoff $Y = 50$.

1. **Risk averse agent.** Let $u(c) = \sqrt{c}$ (a concave utility). Then

$$u(\mathbb{E}[X]) = u(50) \approx 7.07, \quad \mathbb{E}[u(X)] = \frac{1}{2}u(100) + \frac{1}{2}u(0) = \frac{1}{2} \cdot 10 + \frac{1}{2} \cdot 0 = 5.$$

Since $u(\mathbb{E}[X]) > \mathbb{E}[u(X)]$, the agent prefers the sure payoff of 50.

2. **Risk neutral agent.** Let $u(c) = c$ (linear utility). Then

$$u(\mathbb{E}[X]) = u(50) = 50, \quad \mathbb{E}[u(X)] = \frac{1}{2} \cdot 100 + \frac{1}{2} \cdot 0 = 50.$$

The agent is indifferent between the lottery and the sure payoff.

3. **Risk seeking agent.** Let $u(c) = c^2$ (a convex utility). Then

$$u(\mathbb{E}[X]) = u(50) = 2500, \quad \mathbb{E}[u(X)] = \frac{1}{2} \cdot 100^2 + \frac{1}{2} \cdot 0^2 = 5000.$$

Since $u(\mathbb{E}[X]) < \mathbb{E}[u(X)]$, the agent prefers the risky lottery to the certain payoff.

6.2 Pricing a Risky Asset

We describe a simple two-period economy with uncertainty at date $t + 1$.

- The agent receives exogenous endowments e_t and e_{t+1} of the single consumption good at times t and $t + 1$ respectively. These are random variables defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$.
- Preferences are represented by

$$U(c_t, c_{t+1}) = u(c_t) + \beta \mathbb{E}_t[u(c_{t+1})],$$

where $u : \mathbb{R}_+ \rightarrow \mathbb{R}$ is strictly increasing, strictly concave, and continuously differentiable, and $\beta \in (0, 1)$ is the subjective discount factor.

- There is a risky asset that costs p_t units of the good at time t and delivers a random payoff x_{t+1} units of the good at time $t + 1$. If the agent buys $q \in \mathbb{R}$ units of the asset, consumption is

$$c_t = e_t - p_t q, \quad c_{t+1} = e_{t+1} + q x_{t+1}.$$

Theorem 6.7 (Fundamental Pricing Equation). *In the economy described above, any optimal choice q^* satisfies the first-order condition*

$$p_t u'(c_t) = \beta \mathbb{E}_t[u'(c_{t+1}) x_{t+1}].$$

Equivalently, the equilibrium price of the asset is given by

$$p_t = \mathbb{E}_t[m_{t+1} x_{t+1}],$$

for a suitable random variable m_{t+1} .

Proof. The agent chooses q to maximise

$$U(q) := u(e_t - p_t q) + \beta \mathbb{E}_t[u(e_{t+1} + q x_{t+1})].$$

By differentiability of u , the derivative is

$$U'(q) = -p_t u'(e_t - p_t q) + \beta \mathbb{E}_t[u'(e_{t+1} + q x_{t+1}) x_{t+1}].$$

At an optimum q^* we require $U'(q^*) = 0$, so

$$p_t u'(c_t) = \beta \mathbb{E}_t[u'(c_{t+1}) x_{t+1}].$$

Dividing through by $u'(c_t) > 0$ gives

$$p_t = \mathbb{E}_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} x_{t+1} \right],$$

as claimed. □

Remark 6.8 (Interpretation of the factor m in Theorem 6.7). In the utility maximisation setup above, the stochastic discount factor

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

is proportional to the *marginal rate of substitution across time*. It measures how much the agent is willing to give up of certain consumption at date t in exchange for one additional unit of consumption in a particular state at date $t + 1$. A payoff that delivers more consumption precisely in states where m_{t+1} is high is valued more highly, because consumption in those states is especially valuable to the agent.

Remark 6.9 (Generality of the Pricing Equation). The formula $m_{t+1} = \beta u'(c_{t+1})/u'(c_t)$ arises from this specific two-period consumption model. In general financial theory, however, one does not need to assume a particular utility function or even the existence of a representative agent. What matters is the existence of some strictly positive random variable m_{t+1} such that the fundamental relation

$$p_t = \mathbb{E}_t[m_{t+1}x_{t+1}]$$

holds for every traded asset. Any such m is called a *stochastic discount factor*.

Definition 6.10 (Fundamental Pricing Equation). The *fundamental pricing equation* states that for any asset with price p_t at time t and payoff x_{t+1} at time $t + 1$, there exists a stochastic discount factor $m_{t+1} > 0$ such that

$$p_t = \mathbb{E}_t[m_{t+1}x_{t+1}].$$

6.2.1 Examples and Special Cases

Example 6.11 (Riskfree bond). Suppose there is a one-period riskfree bond. Its price at time t is p_t^f , and it pays a deterministic gross return $1 + r$ at time $t + 1$. Thus the payoff is $x_{t+1}^f = 1 + r$.

Applying the fundamental pricing equation gives

$$p_t^f = \mathbb{E}_t[m_{t+1}x_{t+1}^f] = (1 + r) \mathbb{E}_t[m_{t+1}].$$

Normalising the bond to have unit price today ($p_t^f = 1$), we obtain

$$\mathbb{E}_t[m_{t+1}] = \frac{1}{1 + r}.$$

Example 6.12 (Risk-neutral agent). If the agent is risk neutral, then the marginal utility u' is constant. In this case the stochastic discount factor is simply

$$m_{t+1} = \beta \quad (\text{a constant}).$$

From the riskfree bond example we then have

$$\beta = \frac{1}{1+r}.$$

For any asset with payoff x_{t+1} , the pricing equation reduces to

$$p_t = \mathbb{E}_t[\beta x_{t+1}]. \tag{1}$$

where β is given above.

Thus under risk neutrality, the price is the expected payoff discounted at the riskfree rate. This is the pricing equation seen by a risk neutral agent.

Example 6.13 (Stock with dividends). Consider a stock with ex-dividend price p_t at time t , dividend d_{t+1} paid between t and $t+1$, and ex-dividend price p_{t+1} at time $t+1$. The payoff from holding the stock for one period is

$$x_{t+1} = p_{t+1} + d_{t+1}.$$

The fundamental pricing equation then gives

$$p_t = \mathbb{E}_t[m_{t+1}(p_{t+1} + d_{t+1})].$$

This shows that today's price reflects both the expected discounted capital gain (p_{t+1}) and the expected discounted dividend (d_{t+1}). In particular, the presence of dividends simply enters the payoff term, and the same pricing relation continues to hold.

Example 6.14 (European call option). Let $(S_t)_{t \geq 0}$ denote the price process of a stock, with $S_t > 0$ for all t . Fix a strike price $K > 0$ and maturity date $t+1$. The payoff of a European call option written on S is the random variable

$$x_{t+1}^{\text{call}} := (S_{t+1} - K)^+ = \max\{S_{t+1} - K, 0\}.$$

Applying the fundamental pricing equation gives

$$p_t^{\text{call}} = \mathbb{E}_t[m_{t+1} x_{t+1}^{\text{call}}] = \mathbb{E}_t[m_{t+1} (S_{t+1} - K)^+].$$

Thus the option price is the expected discounted payoff, where the nonlinearity of the payoff function $s \mapsto (s - K)^+$ distinguishes options from bonds and stocks. The same pricing relation applies uniformly to this case, only the payoff x_{t+1} changes.

Exercise 6.15 (Lucas exchange economy: deriving the price recursion). Consider a one-good, pure exchange economy in discrete time. The good is perishable. There are n productive units; unit i produces random output Y_{it} at time t , and we write $y_t = (Y_{1t}, \dots, Y_{nt})$. The process (y_t) is Markov with transition density $f(y'|y)$. A representative consumer has preferences

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t U(c_t) \right], \quad 0 < \beta < 1,$$

with $U : \mathbb{R}_+ \rightarrow \mathbb{R}$ strictly increasing and strictly concave.

At each t , let $p_t = (p_{1t}, \dots, p_{nt})$ be the *ex-dividend* price vector. Let q_{it} denote the number of shares of unit i held *after* trading at time t (so $q_{i,t-1}$ are the beginning-of-period holdings at t). One share of unit i pays dividend Y_{it} at t .

Notation. For vectors $a, b \in \mathbb{R}^n$ write $a \cdot b := \sum_{i=1}^n a_i b_i$.

1. **Market clearing.** In a representative-agent equilibrium the net supply of each share is one and all output is consumed:

$$q_{it} = 1 \quad \text{for all } i, t, \quad c_t = \sum_{i=1}^n Y_{it} \quad \text{for all } t.$$

2. **One-period problem at state y .** Fix a current state y and a price function $p(\cdot)$ mapping states to ex-dividend prices. Given beginning-of-period holdings $z \in \mathbb{R}^n$ (these are the $q_{i,t-1}$), the consumer chooses current consumption c and end-of-period holdings $x \in \mathbb{R}^n$ to solve

$$\max_{c \geq 0, x \in \mathbb{R}^n} U(c) + \beta \int v(x, y') f(y'|y) dy'$$

subject to the flow budget constraint

$$c + p(y) \cdot x = \underbrace{y \cdot z}_{\text{dividends}} + \underbrace{p(y) \cdot z}_{\text{value of current portfolio}}.$$

Here $v(x, y')$ is the continuation value at next state y' when future prices follow $p(\cdot)$.

3. **First-order conditions.** Let λ be the multiplier on the flow budget. Show that any solution satisfies

$$U'(c) = \lambda, \quad \lambda p_i(y) = \beta \int \frac{\partial v}{\partial x_i}(x, y') f(y'|y) dy' \quad (i = 1, \dots, n).$$

4. **Marginal value of an extra share.** Suppose the agent carries one additional share of unit i into period $t + 1$, holding everything else fixed. Show that this increases resources in state y' by

$$p_i(y') + Y'_i,$$

since the share pays its dividend Y'_i and can then be sold for its ex-dividend price $p_i(y')$. Deduce that the marginal increase in utility is

$$\frac{\partial v}{\partial x_i}(z, y') = U'(c'(y')) (p_i(y') + Y'_i).$$

5. **Price recursion (general form).** From the first-order and envelope conditions we have, for each i ,

$$p_i(y) = \int \beta \frac{U'(c'(y'))}{U'(c(y))} (p_i(y') + Y'_i) f(y'|y) dy'.$$

Define the stochastic discount factor

$$m(y, y') := \beta \frac{U'(c'(y'))}{U'(c(y))}.$$

Then the recursion becomes

$$p_i(y) = \mathbb{E}[m(y, y') (p_i(y') + Y'_i) \mid y],$$

which is exactly the fundamental pricing equation $p = \mathbb{E}[mx]$ with $x'_i = p_i(y') + Y'_i$ the payoff of share i next period.

Equilibrium form. Imposing market clearing $c(y) = \sum_{j=1}^n y_j$ and $c'(y') = \sum_{j=1}^n y'_j$ gives

$$p_i(y) = \mathbb{E} \left[\beta \frac{U'(\sum_{j=1}^n y'_j)}{U'(\sum_{j=1}^n y_j)} (p_i(y') + Y'_i) \mid y \right],$$

the Lucas pricing recursion.

6. **Risk-neutral special case.** In the special case of risk-neutral agents, we can work out the solution to the recursion. If U is linear, $U'(c)$ is constant. Show that the recursion reduces to

$$p_i(y) = \beta \mathbb{E}[p_i(y') + Y'_i \mid y],$$

and hence

$$p_i(y) = \sum_{s=1}^{\infty} \beta^s \mathbb{E}[Y_{i,t+s} \mid y_t = y],$$

the discounted expected present value of dividends.

Remark 6.16 (Transversality condition). In infinite-horizon models the Euler equations are not enough: we also need to rule out paths where the consumer accumulates wealth or debt without bound. This is expressed by the transversality condition

$$\lim_{T \rightarrow \infty} \mathbb{E}[\beta^T U'(c_T) p_i(y_T) q_{iT}] = 0 \quad \text{for each asset } i.$$

In words: as time goes to infinity, the expected marginal utility value of the portfolio position in any asset must vanish. Otherwise the consumer could improve utility by shifting resources arbitrarily far into the future⁴. The condition simply ensures that the present value of wealth does not explode, so the optimisation problem remains well defined.

Lucas's model showed how asset prices can be derived in general equilibrium when consumption is stochastic and agents are risk averse. The key insight is that prices must satisfy the same fundamental pricing relation $p = \mathbb{E}[mx]$, with the stochastic discount factor $m_{t+1} = \beta U'(c_{t+1})/U'(c_t)$ tied to marginal utility growth. In the special case of risk neutrality, prices reduce to discounted expected present values and discounted prices are martingales. With risk aversion, however, raw prices need not follow martingales; instead, it is the marginal-utility-weighted payoffs that exhibit the martingale property. This clarified why simple martingale models of stock prices are insufficient, and why risk and preferences must enter any coherent theory of asset pricing.

6.3 Equivalent Measures and Risk-Neutral Measures

Up to now, we have written the pricing equation in terms of the stochastic discount factor m_{t+1} :

$$p_t = \mathbb{E}_t[m_{t+1}x_{t+1}].$$

It is often convenient to rewrite this in terms of a change of probability measure.

Equivalent measures

Let Ω be a finite sample space with reference probability \mathbb{P} . A second probability measure \mathbb{Q} on Ω is called *equivalent* to \mathbb{P} if they assign zero probability to exactly the same events:

$$\mathbb{Q}(\omega) = 0 \iff \mathbb{P}(\omega) = 0 \quad \text{for all } \omega \in \Omega.$$

Example 6.17. Suppose $\Omega = \{\omega_1, \omega_2\}$ with $\mathbb{P}(\omega_1) = 0.3$, $\mathbb{P}(\omega_2) = 0.7$. If we define $\mathbb{Q}(\omega_1) = 0.6$, $\mathbb{Q}(\omega_2) = 0.4$, then \mathbb{Q} is equivalent to \mathbb{P} (both assign positive probability to both states). But if $\mathbb{Q}(\omega_1) = 1$, $\mathbb{Q}(\omega_2) = 0$, then \mathbb{Q} is not equivalent, because \mathbb{Q} rules out an event that \mathbb{P} considers possible.

6.3.1 Interpretation of equivalent measures

Why do we introduce a new probability measure \mathbb{Q} when we already have the physical probabilities \mathbb{P} ? The reason is that in markets, payoffs are not valued using physical likelihoods alone. Agents are risk averse, so they attach extra weight to outcomes where consumption is scarce and marginal utility is high. The stochastic discount factor m exactly encodes these adjustments, and defining

⁴Usually this shifting strategy is called a *Ponzi Scheme* and the transversality condition is called a “No Ponzi scheme condition.”

\mathbb{Q} is a convenient way to incorporate them: under \mathbb{Q} , asset prices look like simple discounted expectations.

In practice, this means that an event which is very painful in terms of consumption (say, a recession) receives a higher weight under \mathbb{Q} than under \mathbb{P} . Conversely, an event that occurs in “good times,” when consumption is plentiful and marginal utility is low, is downweighted under \mathbb{Q} . Thus \mathbb{Q} is sometimes called the *risk-adjusted* or *pricing measure*: it reflects how markets collectively price states, not how often those states physically occur.

Example 6.18 (Risk adjustment and equivalent measure in a two-state Lucas economy). Two dates t and $t + 1$. At $t + 1$ there are two states H and T with physical probabilities

$$\mathbb{P}(H) = \mathbb{P}(T) = \frac{1}{2}.$$

There is one perishable good and complete one-period markets. Take the riskfree rate $r = 0$ (so discount factor is 1 for simplicity).

Endowment and preferences. Aggregate endowment at $t + 1$ is

$$C(H) = 4, \quad C(T) = 1,$$

so the economy is “good” in H and “bad” in T . Assume that the preferences are represented by a (representative) agent with log utility,

$$U(c) = \log c.$$

State prices and the risk-adjusted measure. With log utility, the one-period stochastic discount factor is proportional to inverse aggregate consumption:

$$m(H) \propto \frac{1}{C(H)}, \quad m(T) \propto \frac{1}{C(T)}.$$

State prices are

$$\pi(s) = \beta \mathbb{P}(s) m(s), \quad s \in \{H, T\}.$$

With $\beta = 1$ and the normalization $\sum_s \pi(s) = \frac{1}{1+r} = 1$, we get

$$\pi(H) \propto \frac{1}{2} \cdot \frac{1}{4} = \frac{1}{8}, \quad \pi(T) \propto \frac{1}{2} \cdot 1 = \frac{1}{2},$$

so scaling by k so that $\pi(H) + \pi(T) = 1$ (note $\frac{1}{8} + \frac{1}{2} = \frac{5}{8}$) gives $k = \frac{8}{5}$ and hence

$$\pi(H) = \frac{8}{5} \cdot \frac{1}{8} = \frac{1}{5} = 0.2, \quad \pi(T) = \frac{8}{5} \cdot \frac{1}{2} = \frac{4}{5} = 0.8.$$

Because $r = 0$, the *risk-neutral (equivalent) measure* coincides with the normalized state prices:

$$\mathbb{Q}(H) = 0.2, \quad \mathbb{Q}(T) = 0.8.$$

Observe that \mathbb{Q} and \mathbb{P} are equivalent (same null sets), but \mathbb{Q} puts *more* weight on the bad state T and *less* on the good state H . This is the risk adjustment.

Now a natural question is how do we price a physically fair gamble. So consider a wager that pays \$2 in H and \$0 in T . Under \mathbb{P} its expected payoff is \$1, so \$1 looks “fair” under physical probabilities. In Lucas economy equilibrium, the price is the *risk-adjusted* expectation (here $r = 0$ so no discount):

$$p = \mathbb{E}^{\mathbb{Q}}[x] = 2 \cdot \mathbb{Q}(H) + 0 \cdot \mathbb{Q}(T) = 2 \cdot 0.2 = 0.4.$$

Thus the market’s fair price is \$0.40, not \$1. The high-state payoff is downweighted because consumption is plentiful in H (low marginal utility), and upweighted in T where consumption is scarce (high marginal utility).

In the context of asset pricing, equivalent measures formalize the idea that markets weight states by *risk*, not just by *frequency*. Here \mathbb{Q} shifts probability mass toward the low-consumption state. “Fair under \mathbb{P} ” need not be fair in prices; the correct benchmark is \mathbb{Q} .

6.3.2 Equivalent and risk-neutral measures in asset pricing

Example 6.12 notes the general formula for asset pricing seen by a risk neutral agent. So we are motivated to define a risk neutral measure as follows:

Definition 6.19 (Risk-neutral measure). A probability measure \mathbb{Q} equivalent to \mathbb{P} is called *risk-neutral* if, under \mathbb{Q} , every asset price equals the discounted expected payoff:

$$p_t = \frac{1}{1+r} \mathbb{E}_t^{\mathbb{Q}}[x_{t+1}].$$

The next theorem is important because it separates the economics of preferences from the mathematics of pricing. In general, asset prices reflect two ingredients:

- The underlying economics of the environment, such as utility functions of agents and the stochastic process of dividends or endowments. These determine the stochastic discount factor.
- The payoff structure of the asset itself. For example, a bond has a simple payoff, while an option has a more complicated payoff that depends on the price of an underlying stock (See Example 6.14).

The theorem below shows that the economic content of the stochastic discount factor can be absorbed into a new probability measure. Once this change of measure is made, theorems of finance can be developed in a risk neutral world, where asset prices are just discounted expectations of payoffs. This separation allows us to study pricing results without carrying the utility and endowment structure explicitly.

Theorem 6.20 (Risk-neutral pricing). *Let x_{t+1} be the payoff of an asset at $t + 1$ with price p_t at time t . Suppose the pricing equation*

$$p_t = \mathbb{E}_t[m_{t+1}x_{t+1}]$$

holds for some strictly positive stochastic discount factor m_{t+1} . Then there exists a risk neutral probability measure \mathbb{Q} equivalent to \mathbb{P} . In other words,

$$p_t = \frac{1}{1+r} \mathbb{E}_t^{\mathbb{Q}}[x_{t+1}],$$

where $1 + r$ is the gross riskfree return.

Proof. Define \mathbb{Q} by

$$\mathbb{Q}(\omega) = \frac{m_{t+1}(\omega)}{\mathbb{E}_t[m_{t+1}]} \mathbb{P}(\omega).$$

This is a probability measure equivalent to \mathbb{P} , since $m_{t+1} > 0$ and the normalisation ensures $\sum_{\omega} \mathbb{Q}(\omega) = 1$. Then

$$p_t = \mathbb{E}_t[m_{t+1}] \mathbb{E}_t^{\mathbb{Q}}[x_{t+1}].$$

For the riskfree asset with payoff $1 + r$, the pricing equation gives

$$1 = (1 + r) \mathbb{E}_t[m_{t+1}],$$

so $\mathbb{E}_t[m_{t+1}] = 1/(1 + r)$. Substituting yields

$$p_t = \frac{1}{1+r} \mathbb{E}_t^{\mathbb{Q}}[x_{t+1}],$$

as claimed. □

Example 6.21 (Risk-neutral pricing in a two-state model). Fix a time t . Suppose there are two possible states at $t + 1$, called H (“high”) and L (“low”).

- The stock price at time t is denoted $S_t = s$, a positive real number.
- At time $t + 1$ the stock price will be either

$$S_{t+1}(H) = us, \quad S_{t+1}(L) = ds,$$

where $u > 1$ and $0 < d < 1$ are given constants.

- The riskfree gross return is $R = 1 + r > 0$, so one unit invested in the bond at t yields R units at $t + 1$.
- We write x_{t+1} for the payoff of an asset at $t + 1$.

Risk-neutral measure. The risk-neutral pricing theorem states that there exists a probability measure \mathbb{Q} equivalent to the physical measure \mathbb{P} such that

$$p_t = \frac{1}{R} \mathbb{E}_t^{\mathbb{Q}}[x_{t+1}].$$

For the stock itself, the payoff is $x_{t+1} = S_{t+1}$. The condition becomes

$$s = \frac{1}{R} \left(\mathbb{Q}(H) us + \mathbb{Q}(L) ds \right).$$

Computation. Dividing through by $s > 0$, we find

$$1 = \frac{1}{R} (qu + (1 - q)d),$$

where $q := \mathbb{Q}(H)$ and $\mathbb{Q}(L) = 1 - q$. Hence

$$q = \frac{R - d}{u - d}, \quad 1 - q = \frac{u - R}{u - d}.$$

Stock price recursion. Once \mathbb{Q} is determined, the stock price at time t satisfies

$$S_t = \frac{1}{R} \mathbb{E}_t^{\mathbb{Q}}[S_{t+1}] = \frac{1}{R} (q us + (1 - q) ds).$$

This matches the starting condition and confirms that the risk-neutral measure is consistent with stock and bond prices.

Remark 6.22. This example shows how the risk-neutral measure \mathbb{Q} absorbs all the economic content of preferences and endowments. Once \mathbb{Q} is known, the price of any payoff can be computed as a discounted expectation. More complicated securities, such as options, can be handled in exactly the same way by changing only the payoff x_{t+1} .

6.4 Fundamental Theorem of Asset Pricing

We have to get comfortable with convex geometry language in order to locate probability mass functions with required properties. We will explain at the end of the section.

6.4.1 Convex Geometry Preliminaries

Definition 6.23 (Convex Combination). Given points x_1, \dots, x_n in \mathbb{R}^d , a *convex combination* is a linear combination

$$\lambda_1 x_1 + \dots + \lambda_n x_n$$

with coefficients $\lambda_i \geq 0$ and $\sum_{i=1}^n \lambda_i = 1$.

Definition 6.24 (Convex Set). A set $C \subset \mathbb{R}^d$ is *convex* if for all $x, y \in C$ and $\lambda \in [0, 1]$,

$$\lambda x + (1 - \lambda)y \in C.$$

Definition 6.25 (Convex hull). Given a set $A \subset \mathbb{R}^d$, the *convex hull* of A , denoted $\text{conv}(A)$, is the smallest convex set containing A . Equivalently, it is the set of all convex combinations of finitely many points of A :

$$\text{conv}(A) = \left\{ \sum_{i=1}^n \lambda_i x_i : n \in \mathbb{N}, x_i \in A, \lambda_i \geq 0, \sum_{i=1}^n \lambda_i = 1 \right\}.$$

Example 6.26. Convex hulls in the plane

1. If $A = \{x_1, \dots, x_m\} \subset \mathbb{R}^2$ is a finite set of points, then $\text{conv}(A)$ is the polygon with vertices among the x_i .
2. If A is the unit circle in \mathbb{R}^2 , then $\text{conv}(A)$ is the closed unit disc.

These examples show that convex hulls “fill in” the shape generated by a set of points.

The following theorem is key in proving FTAP.

Theorem 6.27 (Hyperplane Separation). *Let $C \subset \mathbb{R}^d$ be a nonempty convex set and let $x \notin C$. Then there exists a nonzero vector $y \in \mathbb{R}^d$ such that*

$$y \cdot x < 0 \quad \text{and} \quad y \cdot z \geq 0 \quad \text{for all } z \in C.$$

6.4.2 Trading Setup

Fix a finite filtered space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t=0}^T, \mathbb{P})$. Each \mathcal{F}_t is a finite partition of Ω into atoms. There are d assets with discounted price processes $\hat{S}^i = (\hat{S}_t^i)_{t=0}^T$, adapted and real-valued. Short sales are allowed.

A quick note about the atoms of the sigma algebras: For an atom $A \in \mathcal{F}_t$, the vector

$$\hat{s}_t(A) := (\hat{S}_t^1(\omega), \dots, \hat{S}_t^d(\omega)) \quad \text{is constant on } A,$$

so we may treat $\hat{s}_t(A) \in \mathbb{R}^d$ unambiguously.

If $A \in \mathcal{F}_{t-1}$ is an atom of \mathcal{F}_{t-1} , then \mathcal{F}_t refines \mathcal{F}_{t-1} , so A can be written as a disjoint union of atoms of \mathcal{F}_t . Denote these atoms by A_1, \dots, A_m , so that

$$A = A_1 \cup \dots \cup A_m, \quad A_j \in \mathcal{F}_t, \quad A_j \subseteq A.$$

We call A_1, \dots, A_m the *descendants* (or “children”) of A at time t . Define the one-step feasible set

$$C(A) := \text{conv} \{ \hat{s}_t(A_1), \dots, \hat{s}_t(A_m) \} \subset \mathbb{R}^d.$$

Next, we set up the basic objects of discrete-time trading.

Definition 6.28 (Trading strategy). Fix d assets with discounted price processes $\widehat{S}^i = (\widehat{S}_t^i)_{t=0}^T$, adapted to the filtration $(\mathcal{F}_t)_{t=0}^T$.

A *trading strategy* is a sequence

$$\phi = (\phi_t)_{t=1}^T, \quad \phi_t = (\phi_t^1, \dots, \phi_t^d),$$

such that for each $t = 1, \dots, T$,

$$\phi_t : \Omega \rightarrow \mathbb{R}^d$$

is \mathcal{F}_{t-1} -measurable (so the holdings at time t are determined by information available at time $t-1$); We call such a process *previsible*.

The real number $\phi_t^i(\omega)$ is interpreted as the number of units of asset i held during the period $(t-1, t]$ in state ω . Note that short selling is allowed since we have allowed negative numbers.

Definition 6.29 (Discounted portfolio value). Let $\phi = (\phi_t)_{t=1}^T$ be a trading strategy, where $\phi_t = (\phi_t^1, \dots, \phi_t^d)$ and each ϕ_t^i is \mathcal{F}_{t-1} -measurable. The associated *discounted portfolio value process* is the adapted process $(\widehat{V}_t)_{t=0}^T$ defined by

$$\widehat{V}_t(\omega) := \sum_{i=1}^d \phi_t^i(\omega) \widehat{S}_t^i(\omega), \quad t = 1, \dots, T,$$

with initial value $\widehat{V}_0 \in \mathbb{R}$ given.

Remark 6.30. Mathematically, \widehat{V}_t is a random variable on Ω , measurable with respect to \mathcal{F}_t . Financially, it represents the value at time t of a portfolio that holds ϕ_t^i units of each asset i . The choice of ϕ_t depends only on information up to $t-1$, while the valuation uses the asset prices \widehat{S}_t^i realized at time t .

Definition 6.31 (Self-financing strategy). Fix d discounted asset price processes

$$\widehat{S}^i = (\widehat{S}_t^i)_{t=0}^T, \quad i = 1, \dots, d,$$

adapted to the filtration $(\mathcal{F}_t)_{t=0}^T$.

Let $\phi = (\phi_t)_{t=1}^T$, where $\phi_t = (\phi_t^1, \dots, \phi_t^d)$ is a trading strategy. The associated discounted portfolio value is

$$\widehat{V}_t := \sum_{i=1}^d \phi_t^i \widehat{S}_t^i, \quad t = 0, 1, \dots, T.$$

We say that ϕ is *self-financing* if

$$\widehat{V}_t - \widehat{V}_{t-1} = \sum_{i=1}^d \phi_t^i (\widehat{S}_t^i - \widehat{S}_{t-1}^i), \quad t = 1, \dots, T.$$

Remark 6.32. The self-financing condition means that changes in portfolio value come entirely from changes in asset prices. In financial terms, the agent never adds or withdraws money from outside the system: wealth evolves only by gains and losses on existing positions.

In any financial model we study, we reject the possibility of *arbitrage*. Intuitively, an arbitrage is a trading strategy that starts with no initial cost, never loses money in any state of the world, and yet produces a strictly positive payoff with positive probability. In other words, it is a “free lunch”: something for nothing. We assume such opportunities cannot persist in competitive markets, and their exclusion is the cornerstone assumption behind the Fundamental Theorem of Asset Pricing. We need to make a precise definition of it now.

Definition 6.33 (Local arbitrage). Fix $t \in \{1, \dots, T\}$ and an atom $A \in \mathcal{F}_{t-1}$. A vector $y \in \mathbb{R}^d$ is called a *local arbitrage on A* if

$$y \cdot \hat{S}_{t-1}(\omega) < 0 \quad \text{for all } \omega \in A,$$

and

$$y \cdot \hat{S}_t(\omega) \geq 0 \quad \text{for all } \omega \in A,$$

with strict inequality for at least one $\omega \in A$.

Remark 6.34. Mathematically, this means: given the information available at time $t - 1$ (represented by the atom A), there exists a portfolio y whose discounted cost is strictly negative at $t - 1$, while its payoff at t is almost surely nonnegative and strictly positive in at least one descendant of A .

Financially, a local arbitrage is an arbitrage opportunity that can be recognised and executed based only on the information available at time $t - 1$. It is the one-step building block of a global arbitrage strategy.

6.4.3 The Fundamental Theorem

In this section, we prove the theorem in a series of lemmas.

Lemma 6.35 (No local arbitrage \Rightarrow convex inclusion). Fix $t \in \{1, \dots, T\}$ and $A \in \mathcal{F}_{t-1}$. If there is no one-step arbitrage on A , then

$$\hat{s}_{t-1}(A) \in C(A).$$

Proof. If $\hat{s}_{t-1}(A) \notin C(A)$, the hyperplane separation lemma (homogeneous form) gives $y \neq 0$ with

$$y \cdot \hat{s}_{t-1}(A) < 0 \quad \text{and} \quad y \cdot z \geq 0 \quad \text{for all } z \in C(A).$$

In particular $y \cdot \hat{s}_t(A_j) \geq 0$ for all children A_j , with strict $>$ for some child because separation is strict. This y is a one-step arbitrage on A , a contradiction. \square

Lemma 6.36. *If $\widehat{s}_{t-1}(A) \in C(A)$, then there exist numbers $q(A \rightarrow A_j) \geq 0$ with $\sum_{j=1}^m q(A \rightarrow A_j) = 1$ such that*

$$\widehat{s}_{t-1}(A) = \sum_{j=1}^m q(A \rightarrow A_j) \widehat{s}_t(A_j).$$

Proof. This is the definition of a convex combination in a convex hull. The coefficients are the convex combination coefficients of $\widehat{s}_{t-1}(A)$ with respect to the vertices $\widehat{s}_t(A_j)$. \square

Lemma 6.37 (Consistent path measure). *Assume Lemma 6.35 holds for every t and every $A \in \mathcal{F}_{t-1}$. For each such A , choose coefficients $q(A \rightarrow A_j) \geq 0$ with $\sum_{j=1}^m q(A \rightarrow A_j) = 1$ such that*

$$\widehat{s}_{t-1}(A) = \sum_{j=1}^m q(A \rightarrow A_j) \widehat{s}_t(A_j),$$

where A_1, \dots, A_m are the descendants of A . Define \mathbb{Q} on atoms $A^{(T)} \in \mathcal{F}_T$ by

$$\mathbb{Q}(A^{(T)}) := \prod_{t=1}^T q\left(A^{(t-1)} \rightarrow A^{(t)}\right),$$

where $A^{(t)}$ is the unique ancestor of $A^{(T)}$ in \mathcal{F}_t . Then \mathbb{Q} is a probability measure on Ω equivalent to \mathbb{P} .

Proof. By construction each $q(A \rightarrow A_j) \geq 0$ and the coefficients from any A sum to one. Thus the product measure \mathbb{Q} assigns nonnegative weights to all terminal atoms, and the total mass is one. If all coefficients are strictly positive, then $\mathbb{Q}(\omega) > 0$ for every $\omega \in \Omega$, hence $\mathbb{Q} \sim \mathbb{P}$ since $\mathbb{P}(\omega) > 0$ as well. If some coefficients vanish, \mathbb{Q} may give zero weight to certain atoms; this can be avoided by choosing convex combinations that place strictly positive weight on each descendant whenever possible (which holds whenever $\widehat{s}_{t-1}(A)$ lies in the relative interior of $C(A)$). In finite models this adjustment is always possible, and does not affect the argument. \square

Lemma 6.38 (Discounted prices are \mathbb{Q} -martingales). *With \mathbb{Q} as in Lemma 6.37, for each asset i and each $t \geq 1$,*

$$\mathbb{E}_{\mathbb{Q}}[\widehat{S}_t^i \mid \mathcal{F}_{t-1}] = \widehat{S}_{t-1}^i.$$

Proof. Fix $A \in \mathcal{F}_{t-1}$. By construction of \mathbb{Q} ,

$$\mathbb{E}_{\mathbb{Q}}[\widehat{S}_t^i \mid A] = \sum_{j=1}^m q(A \rightarrow A_j) \widehat{S}_t^i(A_j).$$

Stacking over $i = 1, \dots, d$ and using Lemma 6.36, the vector of conditional expectations equals $\widehat{s}_{t-1}(A)$, i.e. \widehat{S}_{t-1}^i on A . \square

6.4.4 The theorem

Finally we are ready to prove the fundamental theorem of asset pricing.

Theorem 6.39 (FTAP in finite discrete time). *The following are equivalent:*

1. No arbitrage: *there is no self-financing strategy ϕ with $\widehat{V}_0 = 0$, $\widehat{V}_T \geq 0$ \mathbb{P} -a.s., and $\mathbb{P}(\widehat{V}_T > 0) > 0$.*
2. Equivalent martingale measure: *there exists a probability measure $\mathbb{Q} \sim \mathbb{P}$ such that each discounted price process $\widehat{S}^i = (\widehat{S}_t^i)_{t=0}^T$ is a \mathbb{Q} -martingale.*

Proof. (i) \Rightarrow (ii). Assume (i). Fix $t \in \{1, \dots, T\}$ and an atom $A \in \mathcal{F}_{t-1}$. If $\widehat{s}_{t-1}(A) \notin C(A)$, then by the hyperplane separation lemma there exists $y \in \mathbb{R}^d \setminus \{0\}$ with

$$y \cdot \widehat{s}_{t-1}(A) < 0 \quad \text{and} \quad y \cdot z \geq 0 \quad \text{for all } z \in C(A).$$

In particular $y \cdot \widehat{s}_t(A_j) \geq 0$ for every descendant A_j , with strict $>$ for at least one, so y is a one-step arbitrage on A (by the definition of local arbitrage). Pasting such a one-step arbitrage along the tree yields a global arbitrage, contradicting (i). Hence for every t and $A \in \mathcal{F}_{t-1}$ we must have

$$\widehat{s}_{t-1}(A) \in C(A).$$

This is Lemma 6.35.

For each t and $A \in \mathcal{F}_{t-1}$ choose coefficients of a convex combination $q(A \rightarrow A_j) \geq 0$ with $\sum_j q(A \rightarrow A_j) = 1$ such that

$$\widehat{s}_{t-1}(A) = \sum_j q(A \rightarrow A_j) \widehat{s}_t(A_j),$$

as in Lemma 6.36. Define \mathbb{Q} on terminal atoms by the product

$$\mathbb{Q}(A^{(T)}) := \prod_{t=1}^T q(A^{(t-1)} \rightarrow A^{(t)}),$$

where $A^{(t)}$ is the ancestor of $A^{(T)}$ in \mathcal{F}_t . By Lemma 6.37, \mathbb{Q} is a probability measure on Ω and $\mathbb{Q} \sim \mathbb{P}$. Finally, by Lemma 6.38, for each asset i and each $t \geq 1$,

$$\mathbb{E}_{\mathbb{Q}}[\widehat{S}_t^i \mid \mathcal{F}_{t-1}] = \widehat{S}_{t-1}^i,$$

so each discounted price process is a \mathbb{Q} -martingale. This proves (ii).

(ii) \Rightarrow (i). Assume (ii). Let $\phi = (\phi_t)_{t=1}^T$ be a self-financing strategy with $\widehat{V}_0 = 0$ and discounted value process $(\widehat{V}_t)_{t=0}^T$. By the self-financing identity,

$$\widehat{V}_t - \widehat{V}_{t-1} = \sum_{i=1}^d \phi_t^i (\widehat{S}_t^i - \widehat{S}_{t-1}^i), \quad t = 1, \dots, T.$$

Taking \mathbb{Q} -conditional expectations and using that ϕ_t is \mathcal{F}_{t-1} -measurable (predictable) while each

\widehat{S}^i is a \mathbb{Q} -martingale, we obtain

$$\mathbb{E}_{\mathbb{Q}}[\widehat{V}_t \mid \mathcal{F}_{t-1}] = \widehat{V}_{t-1} + \sum_{i=1}^d \phi_t^i \mathbb{E}_{\mathbb{Q}}[\widehat{S}_t^i - \widehat{S}_{t-1}^i \mid \mathcal{F}_{t-1}] = \widehat{V}_{t-1}.$$

Thus (\widehat{V}_t) is a \mathbb{Q} -martingale with $\widehat{V}_0 = 0$, so $\mathbb{E}_{\mathbb{Q}}[\widehat{V}_T] = 0$. If $\widehat{V}_T \geq 0$ \mathbb{P} -a.s. and $\mathbb{P}(\widehat{V}_T > 0) > 0$, then, by $\mathbb{Q} \sim \mathbb{P}$, we also have $\mathbb{Q}(\widehat{V}_T > 0) > 0$, which forces $\mathbb{E}_{\mathbb{Q}}[\widehat{V}_T] > 0$, a contradiction. Hence no such strategy exists and (i) holds. \square

6.4.5 Applications

Example 6.40 (One-period binomial model). A stock has current price $S_0 = 100$. At time 1 it can go up to $S_1 = 120$ or down to $S_1 = 80$. The riskfree rate is $r = 0.05$, so $1 + r = 1.05$.

Absence of arbitrage requires that there exist a probability $q \in (0, 1)$ such that

$$\frac{S_0}{1} = \frac{1}{1.05} (q \cdot 120 + (1 - q) \cdot 80).$$

Solving gives $q = \frac{1.05 \cdot 100 - 80}{120 - 80} = 0.625$.

Thus under the risk-neutral measure \mathbb{Q} ,

$$\mathbb{Q}(S_1 = 120) = 0.625, \quad \mathbb{Q}(S_1 = 80) = 0.375.$$

Any contingent claim X_1 can now be priced by

$$p_0 = \frac{1}{1.05} \mathbb{E}^{\mathbb{Q}}[X_1].$$

Example 6.41 (Two-period binomial tree). Suppose the stock in the previous example continues to evolve by up and down moves of factor $u = 1.2$ and $d = 0.8$. Starting at $S_0 = 100$, the possible paths are

$$S_2 \in \{100u^2, 100ud, 100d^2\} = \{144, 96, 64\}.$$

From the one-period calculation, the risk-neutral probability of an up move is $q = 0.625$. Hence under \mathbb{Q} the path probabilities are

$$\mathbb{Q}(uu) = q^2, \quad \mathbb{Q}(ud) = q(1 - q), \quad \mathbb{Q}(dd) = (1 - q)^2.$$

Thus \mathbb{Q} is defined consistently on the two-period tree. The FTAP tells us that

$$\frac{S_0}{1} = \frac{1}{(1.05)^2} \mathbb{E}^{\mathbb{Q}}[S_2],$$

and similarly for any claim paying X_2 at time 2.

Example 6.42 (Credit-risk bond). A one-period bond costs p_0 today and pays 100 at time 1 if

there is no default, but pays only 50 if default occurs. Under physical probabilities,

$$\mathbb{P}(\text{no default}) = 0.9, \quad \mathbb{P}(\text{default}) = 0.1.$$

If the riskfree rate is $r = 0.05$, then FTAP implies that there is a risk-neutral probability q such that

$$p_0 = \frac{1}{1.05} (q \cdot 100 + (1 - q) \cdot 50).$$

For instance, if $p_0 = 85$, then

$$85 = \frac{1}{1.05} (50 + 50q) \implies q = \frac{85 \cdot 1.05 - 50}{50} = 0.785.$$

Thus under the risk-neutral measure \mathbb{Q} ,

$$\mathbb{Q}(\text{no default}) = 0.785, \quad \mathbb{Q}(\text{default}) = 0.215.$$

Even though the physical probability of default is 0.1, markets price the bond as if default had probability 0.215. The difference reflects risk premia and is exactly what the FTAP encodes.

6.4.6 Numerical Exercises

Exercise 6.43 (Expected utility and risk attitudes). An agent has utility $u(c) = \sqrt{c}$. The agent faces a lottery that pays \$100 with probability 1/2 and \$0 with probability 1/2.

1. Compute the expected payoff of the lottery.
2. Compute the expected utility of the lottery.
3. Compare the expected utility with the utility of the expected payoff. What does this show about the agent's risk profile?

Exercise 6.44 (Risk neutral vs. risk averse). Consider a lottery that pays \$20 with probability 0.25 and \$0 otherwise.

1. Compute the certainty equivalent if $u(c) = c$ (risk neutral).
2. Compute the certainty equivalent if $u(c) = \sqrt{c}$ (risk averse).
3. Compare the results and interpret.

Exercise 6.45 (Fundamental pricing equation). Suppose $u(c) = \log c$, $\beta = 0.95$. An agent has $c_t = 100$ today and tomorrow's consumption is random:

$$c_{t+1} = \begin{cases} 110 & \text{with probability 0.5,} \\ 90 & \text{with probability 0.5.} \end{cases}$$

An asset pays $x_{t+1} = 10$ in both states.

1. Compute the stochastic discount factor m_{t+1} in each state.
2. Use the fundamental pricing equation to compute the asset's price p_t .

Exercise 6.46 (Equivalent measure). Continue with the previous data.

1. Construct the equivalent measure \mathbb{Q} by

$$\mathbb{Q}(\omega) = \frac{m_{t+1}(\omega)}{\mathbb{E}[m_{t+1}]} \mathbb{P}(\omega).$$

2. Verify that \mathbb{Q} is a probability measure.
3. Recompute the asset price as $p_t = \frac{1}{1+r} \mathbb{E}^{\mathbb{Q}}[x_{t+1}]$.

Exercise 6.47 (Two-state stock pricing). Suppose $S_t = 100$. At $t+1$, the stock is either $S_{t+1} = 120$ or $S_{t+1} = 80$. The riskfree gross return is $R = 1.05$.

1. Determine the risk-neutral probabilities $q = \mathbb{Q}(S_{t+1} = 120)$ and $1 - q$.
2. Verify that the stock price satisfies

$$S_t = \frac{1}{R} \mathbb{E}^{\mathbb{Q}}[S_{t+1}].$$

3. Suppose the stock also pays a dividend $d_{t+1} = 5$ in each state. Compute its price using

$$S_t = \frac{1}{R} \mathbb{E}^{\mathbb{Q}}[S_{t+1} + d_{t+1}].$$