

LECTURES 3 AND 4: MARTINGALES

1. INTRODUCTION

In an arbitrage-free market, the share price of any traded asset at time $t = 0$ is the expected value, under an equilibrium measure, of its discounted price at market termination $t = T$. We shall see that the discounted share price at *any* time $t \leq T$ may also be computed by an expectation; however, this expectation is a *conditional* expectation, given the information about the market scenario that is revealed up to time t . Thus, the discounted share price process, as a (random) function of time, has the property that its value at any time t is the conditional expectation of the value at a future time given information available at t . A stochastic process with this property is called a *martingale*.

The theory of martingales (initiated by JOSEPH DOOB, following earlier work of PAUL LÉVY) is one of the central themes of modern probability. In discrete-time finance, the importance of martingales stems largely from the *Optional Stopping Formula*, which (a) places fundamental constraints on (discounted) share price processes, and (b) provides a computational tool of considerable utility.

2. FILTRATIONS OF PROBABILITY SPACES

Filtrations. In a multiperiod market, information about the market scenario is revealed in stages. Some events may be completely determined by the end of the first trading period, others by the end of the second, and others not until the termination of all trading. This suggests the following classification of events: for each $t \leq T$,

$$(1) \quad \mathcal{F}_t = \{\text{all events determined in the first } t \text{ trading periods}\}.$$

The finite sequence $(\mathcal{F}_t)_{0 \leq t \leq T}$ is a *filtration* of the space Ω of market scenarios. In general, a filtration of a set Ω (not necessarily finite) is defined to be a collection \mathcal{F}_t , indexed by a time parameter t (time may be either discrete or continuous), such that

- (a) each \mathcal{F}_t is a σ -algebra of subsets (events) of Ω ; and
- (b) if $s < t$ then $\mathcal{F}_s \subseteq \mathcal{F}_t$.

Adapted Processes. The share prices of assets in a multiperiod market depend on market scenarios, but evolve in such a way that their values at any time t , being observable at time t , do not depend on the unobservable post- t futures of the scenarios. Thus, the price process S_t of a traded asset is *adapted* to the natural filtration $(\mathcal{F}_t)_{0 \leq t \leq T}$ defined by (1). In general, a sequence X_t of random variables is said to be adapted to a filtration $(\mathcal{F}_t)_{0 \leq t \leq T}$ if, for each t , the random variable X_t is \mathcal{F}_t -*measurable*, that is, if all events of the form $\{\omega : X_t(\omega) \in B\}$, where B is a Borel subset of the real numbers \mathbb{R} , are members of the σ -algebra \mathcal{F}_t .

Example: The Two-Period Binary Market. Recall that there are four market scenarios: $\Omega = \{++, +-, -+, --\}$. For each scenario, the first (respectively, second) entry indicates

whether the share price of the asset STOCK increased or decreased in the first (respectively, second) trading period. There are $2^4 = 16$ events (subsets of Ω) in all. The only events that are determined before the first trading period are the trivial events \emptyset, Ω . There are two other events determined by time $t = 1$: these are

$$\begin{aligned} F_+ &= \{++, +- \} & \text{and} \\ F_- &= \{-+, -- \}. \end{aligned}$$

Consequently, the natural filtration (the filtration defined by equation (1)) is

$$\begin{aligned} \mathcal{F}_0 &= \{\emptyset, \Omega\}, \\ \mathcal{F}_1 &= \{\emptyset, \Omega, F_+, F_-\}, \\ \mathcal{F}_2 &= \{\text{all subsets of } \Omega\}. \end{aligned}$$

Example: The T -Period Binary Market. The set Ω_T of market scenarios consists of the 2^T sequences of length T with entries $+, -$. The *partial scenarios* are the sequences $\xi = (\xi_1 \xi_2 \dots \xi_t)$ of lengths $t \leq T$. For each partial scenario $\xi = (\xi_1 \xi_2 \dots \xi_t)$, define the event $F(\xi)$ to be the set of all (full) scenarios $\omega \in \Omega_T$ whose first t entries are $\xi_1 \xi_2 \dots \xi_t$, that is,

$$(2) \quad F(\xi) = \{\omega \in \Omega : \omega_j = \xi_j \forall j = 1, 2, \dots, t\}.$$

If one visualizes Ω as a binary tree of depth T , then the partial scenarios $\xi = (\xi_1 \xi_2 \dots \xi_t)$ correspond to the nonterminal nodes; for each such ξ , the event $F(\xi)$ consists of all terminal nodes ω lying below the nonterminal node ξ (that is, terminal nodes ω such that the unique downward path in the tree from the root to ω passes through ξ). The σ -algebras \mathcal{F}_t in the natural filtration are given by

$$(3) \quad \mathcal{F}_t = \{\text{all finite unions of events } F(\xi), \text{ where } \xi = \xi_1 \xi_2 \dots \xi_t \in \Omega_t\}.$$

Equivalently, \mathcal{F}_t consists of the events whose *indicator functions*¹ depend on ω only through its first t coordinates. Observe that the \mathcal{F}_t -measurable random variables are the random variables that depend on the market scenario only through its first t coordinates. Thus, a sequence $(X_t)_{0 \leq t \leq T}$ is *adapted* to the filtration $(\mathcal{F}_t)_{0 \leq t \leq T}$ if, for each t and each market scenario $\omega = \omega_1 \omega_2 \dots \omega_T$, the value of $X_t(\omega)$ depends only on the *partial* scenario $\omega_1 \omega_2 \dots \omega_t$.

Exercise: Show that, for the T -period binary market with the natural filtration, there is a one-to-one correspondence between adapted processes and functions on the nodes of the binary tree that represents the market.

3. SELF-FINANCING PORTFOLIOS

Dynamic Portfolios. Consider a T -period market \mathcal{M} with traded assets A^1, A^2, \dots, A^K . Denote by $S_t^A(\omega)$ the share price of asset A at time t under market scenario ω . Generally, a trader will hold not just a single asset, but a *portfolio* consisting of shares (possibly short) of each of the traded assets A^j ; moreover, the trader may adjust his/her portfolio as time progresses, in a manner that is scenario-dependent. Such a time- and scenario-dependent

¹The *indicator function* of an event G is defined to be the random variable $\mathbf{1}_G$ which takes the value 1 if G occurs, and takes the value 0 if G does not occur.

portfolio is called a *dynamic*, or *dynamically rebalanced*, portfolio. If $\theta_t^A(\omega)$ denotes the number of shares of asset A held in a dynamically rebalanced portfolio during the t th trading period (that is, during the period following completion of trading at time t until the beginning of trading at time $t+1$) under scenario ω , then *the sequence $(\theta_t^A)_{0 \leq t \leq T}$ must be adapted to the natural filtration.*² A dynamic portfolio will be called *bounded* if each of the random variables θ_t^A is bounded. Note that in markets with only finitely many scenarios, all dynamic portfolios are bounded.

Self-Financing Portfolios. The total value of the portfolio θ after rebalancing at time t , in scenario ω , is

$$(4) \quad V_t^\theta = V_t^\theta(\omega) = \sum_A \theta_t^A(\omega) S_t^A(\omega),$$

where the sum is over all assets $A = A^1, A^2, \dots, A^K$. Of course, it need not be the case that $V_t^\theta = V_{t+1}^\theta$, as the share prices of the underlying assets A^j will generally change between times t and $t+1$. However, unless the trader invests (or withdraws) additional resources in his/her portfolio at time $t+1$, the total value of the portfolio just before rebalancing at time $t+1$ must be the same as its value just after, that is,

$$\sum_A \theta_t^A(\omega) S_{t+1}^A(\omega) = \sum_A \theta_{t+1}^A(\omega) S_{t+1}^A(\omega);$$

equivalently,

$$(6) \quad V_{t+1}^\theta(\omega) - V_t^\theta(\omega) = \sum_A \theta_t^A (S_{t+1}^A(\omega) - S_t^A(\omega)).$$

A dynamically rebalanced portfolio satisfying (6) is called *self-financing*, because it requires no investments or withdrawals except at the initial time t .

Exercise: Show that, in a T -period market with only finitely many market scenarios, every self-financing portfolio is a linear combination of *pure trading strategies* (as defined in Lecture 2), and every linear combination of pure trading strategies is a self-financing portfolio.

4. MARTINGALES

Let (Ω, \mathcal{F}, P) be a probability space and $(\mathcal{F}_t)_{0 \leq t \leq T}$ or $(\mathcal{F}_t)_{0 \leq t < \infty}$ a filtration by sub- σ -algebras of \mathcal{F} . An adapted sequence X_t of integrable random variables is defined to be a

- **martingale** if $E(X_{t+1} | \mathcal{F}_t) = X_t \forall t$.
- **submartingale** if $E(X_{t+1} | \mathcal{F}_t) \geq X_t \forall t$.
- **supermartingale** if $E(X_{t+1} | \mathcal{F}_t) \leq X_t \forall t$.

Conservation of Expectation. Readers with backgrounds in physics will, no doubt, have some appreciation for the importance of *conservation laws*. In a certain sense, martingales are the stochastic analog of conservation laws: expectation is “conserved”. To understand this, observe that, by the Tower Property of conditional expectation (see the appendix to

²Clairvoyant and psychic traders are exempt from this restriction. However, most economists do not believe that such exceptional traders affect prices in real markets.

this lecture for a brief resume of the properties of conditional expectation) and the fact that the σ -algebras in a filtration are nested,

$$(7) \quad s < t \quad \implies \quad E(X_t | \mathcal{F}_s) = X_s \text{ a.s.}$$

for any martingale $(X_t)_{t \geq 0}$ relative to a filtration $(\mathcal{F}_t)_{0 \leq t < \infty}$. This is true, in particular, for $s = 0$. Consequently, by another application of the Tower Property,

$$(8) \quad EX_t = EX_0 \quad \forall t \geq 0.$$

Simple Examples

(1) Let X be any integrable random variable. Then the sequence X_t defined by $X_t = E(X | \mathcal{F}_t)$ is a martingale, by the Tower Property of conditional expectation. For finite filtrations $(\mathcal{F}_t)_{0 \leq t \leq T}$ of finite probability spaces, these are the *only* martingales (EXERCISE: Prove this.)

(2) Let Y_1, Y_2, \dots be a sequence of independent, identically distributed random variables such that $EY_t = 0$, and let $(\mathcal{F}_t)_{0 \leq t < \infty}$ be the *natural filtration*.³ Then the sequence $X_t = \sum_{j=1}^t Y_j$ is a martingale. This follows easily from the Independence Law for conditional expectation.

(3) Once again let Y_1, Y_2, \dots be a sequence of independent, identically distributed random variables such that $EY_t = 0$ and $EY_t^2 = \sigma^2 < \infty$, and let $(\mathcal{F}_t)_{0 \leq t < \infty}$ be the natural filtration. Then the sequence

$$(9) \quad \left(\sum_{j=1}^t Y_j \right)^2 - \sigma^2 t$$

is a martingale. This is also easy to check. For an interesting use of the martingale (9), see section 8 below.

(4) Let X_1, X_2, \dots be a sequence of independent, identically distributed Bernoulli- p random variables, and let $S_n = \sum_{j=1}^n X_j$. Note that S_n has the binomial- (n, p) distribution. Define

$$Z_n = \left(\frac{q}{p} \right)^{2S_n - n}.$$

Then Z_1, Z_2, \dots is a martingale relative to the usual filtration. Once again, this is easy to check. The martingale $\{Z_n\}_{n \geq 0}$ is quite useful in certain random walk problems — see, for instance, homework problem 2 below.

(5) Let $\{X_n\}_{n \geq 0}$ be a martingale relative to the filtration $\{\mathcal{F}_n\}_{n \geq 0}$, and let $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ be a convex function such that $E\varphi(X_n) < \infty$ for each $n \geq 0$. Then the sequence $\{Z_n\}_{n \geq 0}$ defined by

$$Z_n = \varphi(X_n)$$

³The *natural filtration* is that where \mathcal{F}_t is the σ -algebra generated by Y_1, Y_2, \dots, Y_t , that is, the smallest σ -algebra containing all events of the form $\{Y_i \in B\}$, where B is a Borel subset of \mathbb{R}

is a *submartingale* relative to the filtration $\{\mathcal{F}_n\}_{n \geq 0}$. This is a consequence of the Jensen inequality and the martingale property of $\{X_n\}_{n \geq 0}$:

$$\begin{aligned} E(Z_{n+1} | \mathcal{F}_n) &= E(\varphi(X_{n+1}) | \mathcal{F}_n) \\ &\geq \varphi(E(X_{n+1} | \mathcal{F}_n)) \\ &= \varphi(X_n) = Z_n \end{aligned}$$

Useful special cases: (a) $\varphi(x) = x^2$, and (b) $\varphi(x) = \exp\{\theta x\}$.

5. DISCOUNTED PRICE PROCESSES AS MARTINGALES

The principal (but not the only) reason that the notion of a martingale is of importance in the theory of finance is that discounted price processes in arbitrage-free markets are martingales. Let \mathcal{M} be a T -period market, Ω the (not necessarily finite) set of market scenarios, and $(\mathcal{F}_t)_{0 \leq t \leq T}$ the natural filtration. Assume that \mathcal{M} has a riskless asset with riskless rate of return r per trading period.⁴

Risk-Neutral Probability Measures. A probability measure P on the σ -algebra \mathcal{F}_T is said to be *risk-neutral*, or an *equilibrium measure*, if for every bounded self-financing portfolio θ the value V_0^θ at time $t = 0$ is the expectation, under P , of the discounted value $V_T^\theta = V_T^\theta(\omega)$ at time $t = T$, that is,

$$(10) \quad V_0^\theta = e^{-rT} E_P V_T^\theta.$$

Theorem 1. *If the market \mathcal{M} admits no arbitrages, and has a riskless asset with rate of return $r = 0$, then, under any equilibrium probability measure, the price process $\{S_t\}_{0 \leq t \leq T}$ of any traded asset is a martingale relative to the natural filtration.*

Corollary 1. *If the market \mathcal{M} admits no arbitrages, and has a riskless asset with rate of return $r > 0$, then, under any equilibrium probability measure, the discounted price process of any traded asset $\{e^{-rt} S_t\}_{0 \leq t \leq T}$ is a martingale relative to the natural filtration.*

Proof. It is left as an exercise to verify that the corollary follows from the theorem. To prove the theorem, we must show that for every time $t = 0, 1, 2, \dots, T - 1$,

$$(11) \quad E(S_{t+1} | \mathcal{F}_t) = S_t.$$

By definition of conditional expectation, this is equivalent to showing that for any event $A \in \mathcal{F}_t$,

$$(12) \quad E(S_{t+1} \mathbf{1}_A) = E(S_t \mathbf{1}_A),$$

where S_t is the share price at time t of a traded asset (call it STOCK). For this, we shall exploit the fact that the value of any self-financing portfolio θ at time $t = 0$ is the expected value, under any equilibrium distribution, of the value at the expiration date $t = T$ (this is, in essence, the definition of an equilibrium distribution). Consider the following self-financing portfolio θ :

Strategy θ : At times $\tau = 0, 1, 2, \dots, t$, hold no position, either short or long, in any asset. At time $\tau = t$, if the event A has occurred, sell S_t shares of the riskless asset BOND short,

⁴that is, in every scenario the share price at time t is e^{rt} , for every time $t = 0, 1, \dots, T$.

and use the proceeds to buy one share of STOCK. Hold this position for one trading period, until time $t + 1$; then, no matter what happens in this period, sell the share of STOCK, put the proceeds S_{t+1} into BOND, and hold the resulting position $(-S_t + S_{t+1})$ BOND until termination of the market at $t = T$. In all partial scenarios where event A has *not* occurred at time t , hold no position, either short or long, in any asset, at any time $\tau = t, \dots, T$.

The strategy θ is self-financing, so its value at $t = 0$ is 0. Consequently, under any equilibrium probability, the expected value of the portfolio at $t = T$ must be 0. But the value at termination of the portfolio is $(-S_t + S_{t+1})\mathbf{1}_A$. Therefore,

$$0 = E(S_t - S_{t+1})\mathbf{1}_A \quad \implies \quad ES_t\mathbf{1}_A = ES_{t+1}\mathbf{1}_A.$$

□

6. MARTINGALE TRANSFORMS

The price process of a self-financing, dynamically rebalanced portfolio in a discrete-time multiperiod market with a risk-neutral probability measure P is a special kind of process, called a *martingale transform*. In general, a dynamically rebalanced portfolio holds, during any trading period, a certain number of shares (possibly zero) of each traded asset A in the market. The number θ_t^A of shares of asset A held between times t and $t+1$ must be a function only of the market information available up to time t ; that is, θ_t^A must be \mathcal{F}_t -measurable, where $(\mathcal{F}_t)_{0 \leq t \leq T}$ is the natural filtration for the market. If the portfolio is self-financing, then the fluctuation in the value of the portfolio between the *end* of trading at time t until the beginning of trading at $t + 1$ is

$$(13) \quad V_{t+1}^\theta - V_t^\theta = \sum_A \theta_t(S_{t+1}^A - S_t^A),$$

where the sum is over all traded assets A and, for each such asset, S_t^A is the share price at time t . Notice that each term is the product of θ_t , which is \mathcal{F}_t -measurable, and $S_{t+1}^A - S_t^A$, the increment in the value of A between times t and $t + 1$. Such sequences are called martingale transforms; they are the discrete-time analogs of *Itô integrals*, of which we will hear more later. To define martingale transforms in general, we must introduce the notion of a *predictable* sequence.

Predictable Sequences. Let $\{\mathcal{F}_t\}_{t \in \mathbb{Z}_+}$ be a discrete-time filtration, and let $\{Y_t\}_{t \in \mathbb{N}}$ be a sequence of random variables adapted to \mathcal{F}_t . The sequence Y_t is said to be *predictable* (relative to the filtration \mathcal{F}_t) if, for every $t \geq 1$, the random variable Y_t is \mathcal{F}_{t-1} -measurable. If \mathcal{F}_t is the natural filtration for a discrete-time, multiperiod market, then the sequence Y_t is predictable if the value at the *end* of each trading period is a function only of the market information available up to the *beginning* of the trading period. For instance, if Y_{t+1} is the number of shares of an asset A held between times t and $t + 1$, then the sequence Y_t must be predictable. (NOTE: It is more conventional to index the number of shares held during the period from t until $t + 1$ by t .)

Martingale Transforms. Let $\{X_t\}_{t \in \mathbb{Z}_+}$ be a martingale (relative to the filtration \mathcal{F}_t), and let $\{Y_t\}_{t \in \mathbb{Z}_+}$ be a predictable sequence. The *martingale transform* $\{(Y \cdot X)_t\}_{t \in \mathbb{Z}_+}$ is the sequence of random variables defined as follows:

$$(14) \quad (Y \cdot X)_t = X_0 + \sum_{j=0}^{t-1} Y_j(\Delta X)_{j+1} \quad \text{where}$$

$$(\Delta X)_{t+1} = X_{t+1} - X_t.$$

The same equations may be used to define *submartingale transforms* and *supermartingale transforms*. If the sequence $\{X_t\}$ is a *sub-* (respectively, *super-*)martingale, then the sequence $(Y \cdot X)_t$ defined by equations (14) is called a sub- (respectively, super-)martingale transform.

NOTE: There are different conventions regarding the initial term of the sequence.

Proposition 1. *Assume that $\{X_t\}_{t \in \mathbb{Z}_+}$ is an adapted sequence and $\{Y_t\}_{t \in \mathbb{Z}_+}$ a predictable sequence of bounded random variables, both relative to a filtration $(\mathcal{F}_t)_{0 \leq t < \infty}$. If $\{X_t\}_{t \in \mathbb{Z}_+}$ is a martingale, then the martingale transform $\{(Y \cdot X)_t\}_{t \in \mathbb{Z}_+}$ is a martingale relative to $(\mathcal{F}_t)_{0 \leq t < \infty}$. If $\{X_t\}_{t \in \mathbb{Z}_+}$ is a sub- (respectively, super-)martingale, and if in addition the random variables Y_t are nonnegative, then the transform sequence $\{(Y \cdot X)_t\}_{t \in \mathbb{Z}}$ is a sub- (respectively, super-)martingale.*

Proof. Homework. □

Proposition 1 leads immediately to the following important extension of Theorem 1.

Theorem 2. *Assume that \mathcal{M} is a multiperiod market containing a riskless asset with rate of return $r = 0$, and assume that there is a risk-neutral probability measure P for \mathcal{M} . Then for every self-financing portfolio θ , the value process $(V_t^\theta)_{0 \leq t \leq T}$ is a martingale under P .*

Proof. By Theorem 1, the share price process S_t^A of each traded asset A is a martingale. If θ is a self-financing portfolio, then, for each asset A , the number θ_{t-1}^A of shares of A held during the period from time $t - 1$ until t is \mathcal{F}_{t-1} -measurable, and so the sequence $(\theta_{t-1}^A)_{0 \leq t \leq T}$ is predictable. Hence, for each asset A , the process

$$V_t^{\theta, A} = \sum_{j=1}^t \theta_{j-1}^A (S_j^A - S_{j-1}^A),$$

which gives the cumulative change in the value of the portfolio due to fluctuations in the share price of A , is a martingale transform. Therefore, by Proposition 1, it is a martingale. By the definition of a self-financing portfolio, V_t^θ is the initial value V_0^θ plus the sum of the values $V_t^{\theta, A}$ over all assets A . Since the sum of martingales is a martingale, it follows that the value process $(V_t^\theta)_{0 \leq t \leq T}$ is a martingale. □

7. OPTIONAL STOPPING

The cornerstone of martingale theory is Doob's *Optional Stopping Theorem*. This states, roughly, that “stopping” a martingale at a random time τ does not alter the expected “payoff”, provided the decision about when to stop is based solely on information available

up to τ . Such random times are called *stopping times*.⁵ The formal definition is as follows: An integer-valued random variable $\tau = \tau(\omega)$ is a *stopping time* relative to a filtration $(\mathcal{F}_t)_{0 \leq t \leq T}$ or $(\mathcal{F}_t)_{0 \leq t < \infty}$ if, for every $n = 0, 1, 2, \dots$, the event that $\tau = n$ is an element of \mathcal{F}_n :

$$(16) \quad \{\omega : \tau(\omega) = n\} \in \mathcal{F}_n.$$

Exercise: (1) Check that this is equivalent to requiring that $\{\omega : \tau(\omega) \leq n\} \in \mathcal{F}_n$ for all n . (2) Show that if τ is a stopping time, then so is $\tau \wedge n$, for any (nonrandom) integer n . (3) Show that if τ and ν are stopping times, then so are $\tau \wedge \nu$ and $\tau \vee \nu$. (The notation \wedge denotes *minimum*, and \vee denotes *maximum*.)

Theorem 3. (*Optional Stopping Theorem*) Let $\{X_n\}_{n \in \mathbb{Z}_+}$ be a martingale, submartingale, or supermartingale relative to a filtration $\{\mathcal{F}_n\}_{n \in \mathbb{Z}_+}$, and let τ be a stopping time. Then for any $n \in \mathbb{N}$,

$$(17) \quad EX_{\tau \wedge n} = EX_0 \quad (\text{martingales})$$

$$(18) \quad EX_{\tau \wedge n} \geq EX_0 \quad (\text{submartingales})$$

$$(19) \quad EX_{\tau \wedge n} \leq EX_0 \quad (\text{supermartingales})$$

Proof. See below for the proof of (17), and see the homework section for the proofs of (18)-(19). □

The direct implication for asset prices is that one cannot get a better expected discounted return on an asset by holding it until a stopping time than one can get by simply holding it until a nonrandom termination date $t = T$. This should not be surprising, because if it were possible to get a better expected return, then one could create an arbitrage (exercise: explain how). But this is not why the Optional Sampling Theorem is really important. The homework problems below give some indication of how more subtle applications of the theorem lead to nontrivial results in asset pricing and random walk theory.

The easiest proof of the Optional Stopping Theorem is based on the fact that martingale transforms are martingales (Proposition 1 above).

Proposition 2. Let τ be a stopping time relative to the filtration \mathcal{F}_t . If $\{X_t\}_{t \geq 0}$ is a martingale, then the sequence of random variables $\{X_{\tau \wedge t}\}_{t \geq 0}$ is a martingale transform, and therefore a martingale. Similarly, if $\{X_t\}_{t \geq 0}$ is a sub- (respectively, super-) martingale, then so is $\{X_{\tau \wedge t}\}_{t \geq 0}$.

Proof. Homework problem. (Identify the appropriate predictable sequence Y_n , prove that it is predictable, and prove that $X_{\tau \wedge n} = (Y \cdot X)_n$.) □

The Optional Stopping Theorem now follows from Conservation of Expectation (8). □

⁵In some of the older literature, they are called *Markov times* or *optional times*.

8. EXAMPLE: THE GAMBLER'S RUIN PROBLEM

The Game. Two gamblers, FATS and SLIM, play the following game: FATS repeatedly tosses a fair coin. After each toss that comes up H , SLIM pays FATS one dollar. After each toss that comes up T , FATS pays SLIM one dollar. The game continues until either one or the other gambler runs out of money. If FATS starts with $\$A$ and SLIM starts with $\$B$,

- (A) What is the probability that, when the game ends, FATS has all the cash?
- (B) What is the expected duration of the game?

We may recast this as an optional stopping problem. Let X_1, X_2, \dots be the sequence of play-by-play increments in FATS' fortune: thus, $X_i = \pm 1$ according to whether the i th toss is H or T . The total change in FATS' fortune after n plays is

$$S_n = \sum_{i=1}^n X_i.$$

The game continues until time τ , where

$$\tau = \min\{n : S_n = +A \text{ or } -B\}.$$

It is not difficult to see that τ is a stopping time relative to the natural filtration $\mathcal{F}_n = \sigma(X_1, X_2, \dots, X_n)$. Moreover, the sequence S_n is a martingale relative to the natural filtration, by example (2) in section 4 above. Consequently, by the Optional Stopping Theorem, for each $n < \infty$,

$$\begin{aligned} 0 = ES_0 &= ES_{\tau \wedge n} = AP\{\tau \leq n \text{ and } S_\tau = A\} \\ &\quad - BP\{\tau \leq n \text{ and } S_\tau = -B\} \\ &\quad + ES_n \mathbf{1}_{\{\tau > n\}}. \end{aligned}$$

As $n \rightarrow \infty$, the probability of the event $\tau > n$ converges to zero. (Explanation: If at any time there are $A+B$ consecutive Heads, or $A+B$ consecutive Tails then the game must end, because one or the other player will be ruined. If n is extremely large, the chance that there has *not* been a run of $A+B$ consecutive Heads in the first n tosses is very small.) Since S_n must be between A and $-B$ on the event $\tau > n$, it follows that the last expectation above converges to zero as $n \rightarrow \infty$. Thus, letting $n \rightarrow \infty$ yields the following simplified equation:

$$0 = AP\{S_\tau = A\} - BP\{S_\tau = -B\}.$$

Since S_τ must be either A or $-B$, the two probabilities in this equation must sum to 1. Hence, we have two equations in two unknowns, which we may solve to obtain the solution to Problem (A):

$$(23) \quad \boxed{P\{S_\tau = A\} = \frac{B}{A+B}.$$

To solve Problem (B), we shall again appeal to the Optional Stopping Theorem, but this time using the martingale $S_n^2 - n$ (see example (3) of section 4 above). The Optional Stopping

Theorem implies that, for each $n = 1, 2, \dots$,

$$(24) \quad \begin{aligned} E(S_{\tau \wedge n}^2 - \tau \wedge n) &= 0 & \implies \\ E\tau \wedge n &= ES_{\tau \wedge n}^2 \\ &= ES_{\tau}^2 \mathbf{1}_{\{\tau \leq n\}} + ES_n^2 \mathbf{1}_{\{\tau > n\}}. \end{aligned}$$

Now $\tau \wedge n \rightarrow \tau$ and $S_{\tau}^2 \mathbf{1}_{\{\tau \leq n\}} \rightarrow S_{\tau}^2$ as $n \rightarrow \infty$, and in both cases the convergence is monotone. Consequently, by the Monotone Convergence Theorem,

$$\begin{aligned} \lim_{n \rightarrow \infty} E\tau \wedge n &= E\tau & \text{and} \\ \lim_{n \rightarrow \infty} ES_{\tau}^2 \mathbf{1}_{\{\tau \leq n\}} &= ES_{\tau}^2 \\ &= A^2 \left(\frac{B}{A+B} \right) + B^2 \left(\frac{A}{A+B} \right) \\ &= AB \end{aligned}$$

the last by the result obtained in solving problem (A) above. Since S_n^2 is bounded on the event $\tau > n$, and since the probability of this event converges to zero as $n \rightarrow \infty$, the last expectation in (24) above converges to zero as $n \rightarrow \infty$. Therefore, as $n \rightarrow \infty$ the equation (24) simplifies to

$$(25) \quad \boxed{E\tau = AB}$$

9. THE MARTINGALE REPRESENTATION THEOREM AND HEDGING

In a market with a risk-neutral probability measure, where the riskless rate of return is $r = 0$, the value process of every bounded, self-financing portfolio is a martingale. Thus, it seems a natural question to ask is this: *How does one characterize the set of all martingales?* To answer this in general is too ambitious a goal; however, in the important special case of the *homogeneous, binary, T-period market* \mathcal{M} , we may give a complete answer.

Recall that the market \mathcal{M} has scenario space $\Omega = \{+, -\}^T$, the set of all sequences of pluses and minuses of length T . There is a riskless asset BOND with rate of return r , and a risky asset STOCK whose price process evolves according to the rule

$$(26) \quad \begin{aligned} S_{t+1}(\omega_1 \omega_2 \dots \omega_t +) &= S_t(\omega_1 \omega_2 \dots \omega_t)u; \\ S_{t+1}(\omega_1 \omega_2 \dots \omega_t -) &= S_t(\omega_1 \omega_2 \dots \omega_t)d. \end{aligned}$$

If the riskless rate of return is $r = 0$ (as is the case if share prices in the market are denominated in shares of BOND), then the risk-neutral probability measure P is the probability measure on Ω under which the *coordinate random variables* ξ_t , defined by

$$(27) \quad \xi_t(\omega_1 \omega_2 \dots \omega_T) = \omega_t \cdot 1 \quad \text{for } t = 1, 2, \dots, T,$$

are independent and identically distributed, with distribution

$$(28) \quad \begin{aligned} P\{\xi_t = +1\} &= p & \text{and} \\ P\{\xi_t = -1\} &= q := 1 - p & \text{where} \\ p &= (1 - d)/(u - d). \end{aligned}$$

The natural filtration $(\mathcal{F}_t)_{0 \leq t \leq T}$ of the scenario space Ω has $\mathcal{F}_t = \sigma(\xi_1, \xi_2, \dots, \xi_t)$, that is, \mathcal{F}_t contains all events whose indicator functions are functions of the coordinate random variables, $\xi_1, \xi_2, \dots, \xi_t$.

Theorem 4. (*Martingale Representation Theorem*) *If $(Y_t)_{0 \leq t \leq T}$ is a martingale relative to the natural filtration $(\mathcal{F}_t)_{0 \leq t \leq T}$, then $(Y_t - Y_0)_{0 \leq t \leq T}$ is a martingale transform of the STOCK price martingale $(S_t)_{0 \leq t \leq T}$; that is, there exists a predictable sequence $(\beta_t)_{1 \leq t \leq T}$ such that for each $t = 1, 2, \dots, T$,*

$$(29) \quad Y_t = Y_0 + \sum_{j=1}^t \beta_j (S_j - S_{j-1}).$$

Why is this of interest in finance? Recall that the share price process of *any* traded asset must be a martingale under the risk-neutral measure P . In particular, if derivative securities (for example, call options) are included in the market, their share price processes must be martingales, and therefore, by the Martingale Representation Theorem, must be of the form (29). The importance of equation (29) is that it provides a self-financing portfolio in the assets STOCK and BOND whose value process is Y_t . Thus, for every derivative security there is a self-financing replicating portfolio in STOCK and BOND. The portfolio is defined using the predictable sequence β_t in the martingale transform (29): in the time period from t until $t+1$, the portfolio holds Y_0 shares of BOND and $\theta_t = \beta_{t+1}$ shares of STOCK. Observe that the process $(\theta_t)_{0 \leq t \leq T-1}$ is adapted to the natural filtration (as it must be) because the sequence β_t is predictable.

Proof of the Martingale Representation Theorem. Fix a scenario $\omega = \omega_1 \omega_2 \dots \omega_T \in \Omega$, and let $G_t(\omega)$ be the set of all (full) scenarios whose first t entries are $\omega_1 \omega_2 \dots \omega_t$. Since the sequence $(Y_t)_{0 \leq t \leq T}$ is a martingale, it must be the case that, for every $t < T$ and every $\omega \in \Omega$,

$$(30) \quad EY_{t+1} \mathbf{1}_{G_t(\omega)} = EY_t \mathbf{1}_{G_t(\omega)}.$$

Moreover, since the sequence $(Y_t)_{0 \leq t \leq T}$ is adapted to the natural filtration, the value $Y_{t+1}(\omega) - Y_t(\omega)$ depends on the scenario $\omega = \omega_1 \omega_2 \dots \omega_T$ only through its first $t+1$ entries. Thus, equation (30) implies that, for each t and each ω ,

$$(31) \quad pY_{t+1}(\omega_1 \omega_2 \dots \omega_{t+}) + qY_{t+1}(\omega_1 \omega_2 \dots \omega_{t-}) = Y_t(\omega_1 \omega_2 \dots \omega_t).$$

This equation also holds with Y_{t+1} and Y_t replaced by S_{t+1} and S_t , respectively, because $(S_t)_{0 \leq t \leq T}$ is also a martingale. Solving both equations for $-q/p$ leads to the relation

$$\frac{Y_{t+1}(\omega_1 \omega_2 \dots \omega_{t+}) - Y_t(\omega_1 \omega_2 \dots \omega_t)}{Y_{t+1}(\omega_1 \omega_2 \dots \omega_{t-}) - Y_t(\omega_1 \omega_2 \dots \omega_t)} = -\frac{q}{p} = \frac{S_{t+1}(\omega_1 \omega_2 \dots \omega_{t+}) - S_t(\omega_1 \omega_2 \dots \omega_t)}{S_{t+1}(\omega_1 \omega_2 \dots \omega_{t-}) - S_t(\omega_1 \omega_2 \dots \omega_t)},$$

which in turn implies that

$$\begin{aligned}
(33) \quad & \frac{Y_{t+1}(\omega_1\omega_2 \dots \omega_{t+}) - Y_t(\omega_1\omega_2 \dots \omega_t)}{S_{t+1}(\omega_1\omega_2 \dots \omega_{t+}) - S_t(\omega_1\omega_2 \dots \omega_t)} \\
& = \frac{Y_{t+1}(\omega_1\omega_2 \dots \omega_{t-}) - Y_t(\omega_1\omega_2 \dots \omega_t)}{S_{t+1}(\omega_1\omega_2 \dots \omega_{t-}) - S_t(\omega_1\omega_2 \dots \omega_t)} \\
& := \beta_t(\omega_1\omega_2 \dots \omega_t).
\end{aligned}$$

Notice that the common value of the fractions on the two sides of (33) depends only on $\omega_1\omega_2 \dots \omega_t$, so the definition of β_t is valid. The relation (29) is now evident.

10. APPENDIX: CONDITIONAL EXPECTATION

Let (Ω, \mathcal{F}, P) be a probability space and let \mathcal{G} be a σ -algebra such that $\mathcal{G} \subset \mathcal{F}$.

Theorem 5. *Let X be any random variable such that $E|X| < \infty$. Then there is a unique⁶ \mathcal{G} -measurable random variable $E(X|\mathcal{G})$ such that, for every bounded \mathcal{G} -measurable random variable Y ,*

$$(34) \quad E(XY) = E(E(X|\mathcal{G})Y).$$

The random variable $E(X|\mathcal{G})$ is called the **conditional expectation** of X given \mathcal{G} . Theorem 5 is one of the major results of measure theory (it is roughly equivalent to the *Radon–Nikodym* theorem for finite measures), and you may find its proof in most X -rated probability texts.

It is often easy to guess the conditional expectation of a random variable X given a σ -algebra \mathcal{G} . Verification of the guess using the definition (34) may be tedious, as it requires considering *all* \mathcal{G} -measurable random variables Y . Fortunately, this is not always necessary: in fact, to prove that $Z = E(X|\mathcal{G})$, it is enough to verify that Z is \mathcal{G} -measurable, and that

$$(35) \quad E(X\mathbf{1}_A) = E(Z\mathbf{1}_A) \quad \forall A \in \mathcal{G}$$

(here $\mathbf{1}_A$ is the *indicator function* of the event A). This follows by a standard approximation argument of measure theory: (1) If (35) holds for all indicator random variables $\mathbf{1}_A$, then $E(XY) = E(ZY)$ for all *simple functions* (linear combinations of indicator functions), by the linearity of expectation. (2) If $E(XY) = E(ZY)$ holds for all simple functions Y , then it holds for all monotone limits Y of simple functions, by the Monotone Convergence Theorem. (3) All bounded \mathcal{G} -measurable random variables are monotone limits of \mathcal{G} -measurable simple functions. For further details, consult your local measure-theorist.

Example: The binary tree. Consider, once again, the T -period binary market, and let $\{\mathcal{F}_t\}$ be the natural filtration. Let P be any probability measure on \mathcal{F}_T assigning positive probability to every scenario:

$$\pi(\omega) := P(\{\omega\}) > 0 \quad \forall \omega \in \Omega.$$

⁶Unique, that is, up to changes on sets of probability zero

For each $t \leq T$ and each scenario $\omega = \omega_1\omega_2 \dots \omega_T$, let $\xi_t = \omega_1\omega_2 \dots \omega_t$ denote the *partial* scenario coinciding with ω up to time t , and define

$$(36) \quad \pi(\omega | \mathcal{F}_t) := \frac{\pi(\omega)}{\sum_{\omega' \in F(\xi_t)} \pi(\omega')}.$$

The sum in the denominator is over all full scenarios ω' coinciding with ω up to time t . Thus, the ratio is the conditional probability (as commonly defined in elementary discrete probability) of the scenario ω , given that ω is not ruled out by what happens in the first t periods. For any random variable X defined on Ω , the conditional expectation $E(X|\mathcal{F}_t)$ is obtained by averaging against the conditional probability distribution $\pi(\omega|\mathcal{F}_t)$, that is,

$$(37) \quad E(X|\mathcal{F}_t)(\omega) = \sum_{\omega' \in F(\xi_t)} X(\omega')\pi(\omega' | \mathcal{F}_t).$$

Exercise: Verify that this is true; that is verify that the function of ω defined by the right side of this equation is the unique \mathcal{F}_t -measurable random variable such that equation (34) holds for every random variable Y .

Properties of Conditional Expectation. One of the advantages of defining conditional expectation via equation (34) is that it makes the following elementary properties transparent.

- (1) **Linearity:** $E(aX + bY|\mathcal{G}) = aE(X|\mathcal{G}) + bE(Y|\mathcal{G})$ for all $a, b \in \mathbb{R}$.
- (2) **Positivity:** If $X \geq 0$ then $E(X|\mathcal{G}) \geq 0$.
- (3) **Stability:** If X is \mathcal{G} -measurable, then $E(XY|\mathcal{G}) = XE(Y|\mathcal{G})$.
- (4) **Independence Law:** If X is independent of \mathcal{G} then $E(X|\mathcal{G}) = EX$ is constant.
- (5) **Tower Property:** If $\mathcal{H} \subset \mathcal{G}$ then $E(X|\mathcal{H}) = E(E(X|\mathcal{G})|\mathcal{H})$.
- (6) **Projection Law:** $E(E(X|\mathcal{G})) = EX$.
- (7) **Jensen Inequalities:** If $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ is convex and $E|X| < \infty$ then

$$E(\varphi(X)) \geq \varphi(EX) \text{ and} \\ E(\varphi(X)|\mathcal{G}) \geq \varphi(E(X|\mathcal{G})).$$

With the exception of the Jensen inequality (7) and the Independence Law (4), all of these properties may be proved easily, using only Theorem 5. You should try to check as many as you can. The Independence Law is not difficult to establish for filtrations of *finite* probability spaces, but proof in general requires some facts from measure theory.

11. PROBLEMS

1. **Bonds.** A (zero-coupon) *bond* with maturity M is a contract that pays the owner \$1 at the maturity date M . Consider a T -period market in which bonds of all maturities $M = 1, 2, \dots, T$ are traded; denote by $B_{t,M} = B_{t,M}(\omega)$ the price (in dollars) at time t of one maturity- M bond.

(a) What is the riskless asset in this market?

(b) A *coupon-bearing* bond is a contract that pays the owner an amount C_t dollars at time t , for $t = 1, 2, \dots, T$. Find a formula for the price at time $t = 0$ of such a contract.

(c) Let π be an equilibrium distribution for the market, and let E_π denote the expectation operator for the distribution π . For each $t = 0, 1, 2, \dots, T$, define $Z_t = \prod_{j=0}^{t-1} B_{j,j+1}$. Prove that the sequence Z_t is a martingale relative to the usual filtration.

2. **Gambler's Ruin, Revisited.** The game is exactly the same as described in section 8 *except* that the coin is unfair: the probability that it comes up H is $p > 1/2$, and the probability that it comes up T is $q = 1 - p < 1/2$. The problem is, once again, to determine the probability that FATS exhausts his initial fortune before SLIM exhausts his.

(a) Define X_i and S_n as in section 8. Show that

$$Z_n := \left(\frac{q}{p}\right)^{S_n}$$

is a martingale relative to the usual filtration.

(b) Use the Optional Stopping Theorem for the martingale Z_n to determine the probability that FATS runs out of money before he bankrupts SLIM.

3. Prove Proposition 1 in the notes.

4. Prove Proposition 2 in the notes.

5. **A Barrier Option:** This contract pays the owner 1 share of STOCK at the first time $t = 1, 2, \dots$ that the share price of STOCK exceeds $\$A$, or at $t = T$ if the share price of STOCK does not exceed $\$A$ at any $t = 1, 2, \dots, T$. Evaluate the market price of this contract at $t = 0$ in two different ways:

(a) by an arbitrage argument; and

(b) using the Optional Sampling Theorem of martingale theory.

6. **American Call Option:** The American call with strike K is a contract that gives the owner the right to purchase one share of STOCK for $\$K$ at *any* time $t = 1, 2, \dots, T$. (The European call may only be exercised at $t = T$.) The owner's decision about when to exercise the option must be based only on information available to him/her at the exercise time — in mathematical terms, the time τ at which the option is exercised must be a *stopping time* relative to the natural filtration.

Assume that there is a riskless asset BOND whose rate of return is $r > 0$. Use the Jensen inequality for conditional expectation to prove that early exercise of the American call is *suboptimal*. Thus, the American call and the European call have the same market value.

7. **American Put Option:** The American *put* with strike K is a contract that gives the owner the right to *sell* one share of STOCK for $\$K$ at *any* time $t = 1, 2, \dots, T$. Assume that there is a riskless asset BOND whose rate of return is $r > 0$. Give an example to show that, in certain markets and in certain circumstances, it may be better to exercise the put option early.

HINT: Consider a homogeneous 2-period binary market. Show that for some values of u, d, r and K , you get a higher expected payoff by exercising the put at $t = 1$ when the partial scenario is — than waiting until $t = 2$. Note that if $r = 0$, then there is no advantage to early exercise.

8. **Doob Decomposition Theorem:** Let $\{X_n\}_{n \geq 0}$ be a submartingale relative to the filtration $\{\mathcal{F}_n\}_{n \geq 0}$. Prove that there is an increasing, predictable sequence $\{Y_n\}_{n \geq 1}$ and a martingale $\{Z_n\}_{n \geq 0}$ such that

$$X_n = Z_n + Y_n.$$

NOTE: The sequence $\{Y_n\}$ is called the *compensator* for the submartingale $\{X_n\}$. The terminology should help you to guess how to define the random variables Z_n and Y_n . Once you guess what the decomposition is, it should be easy to check, using the basic properties of conditional expectation, that the sequences Z_n and Y_n have the desired properties.