s borrowing rate in the UK is u s to set the exchange rate, C_T, in to exchange C_T dollars for one , what is the fair value of C_T?

tion to a speculator. This amounts f is a fixed amount of cash if the if it goes down. If the speculator option writer be willing to write are zero.

ple 1.6.6 operates in the Sterling ng tradables and find the corre- ne as the risk-neutral probabilities cost at time zero of the option as

llar trader uses the dollar bond as ader uses a Sterling bond.

2 Binomial trees and discrete parameter martingales

Summary

In this chapter we build some more sophisticated market models that track the evolution of stock prices over a succession of time periods. Over each individual time period, the market follows our simple binary model of Chapter 1. The possible trajectories of the stock prices are then encoded in a tree. A simple corollary of our work of Chapter 1 will allow us to price claims by taking expectation with respect to certain probabilities on the tree under which the stock price process is a discrete parameter *martingale*.

Definitions and basic properties of discrete parameter martingales are presented and illustrated in §2.3, and we see for the first time how martingale methods can be employed as an elegant computational tool. Then, §2.4 presents some important martingale theorems. In §2.5 we pave the way for the Black–Scholes analysis of Chapter 5 by showing how to construct, in the martingale framework, the portfolio that replicates a claim. In §2.6 we preview the Black–Scholes formula with a heuristic passage to the limit.

2.1 The multiperiod binary model

Our single period binary model is, of course, inadequate as a model of the evolution of an asset price. In particular, we have allowed ourselves to observe the market at just two times, zero and T. Moreover, at time T, we have supposed the stock price to take one of just two possible values. In this section we construct more sophisticated market models by stringing together copies of our single period model into a tree.

Once again our financial market will consist of just two instruments, the stock and a cash bond. As before we assume that unlimited amounts of both can be bought and sold without transaction costs. There is no risk of default on a promise and the market is prepared to buy and sell a security for the same price (that is, there is no *bid–offer spread*).

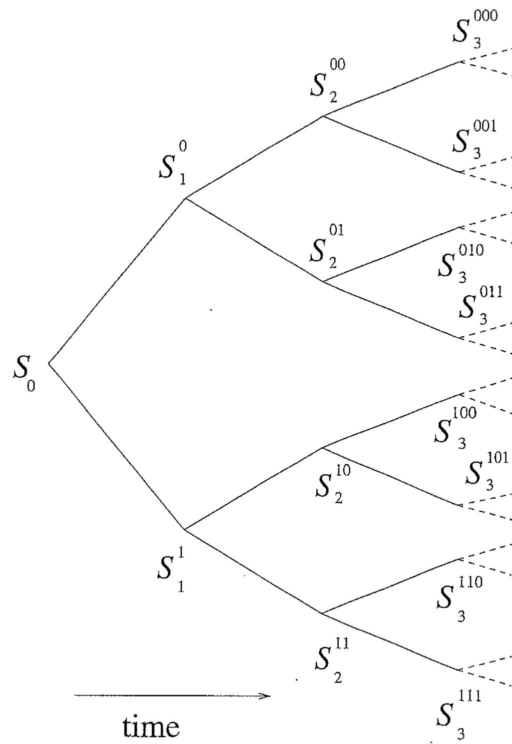

Figure 2.1 The tree of stock prices.

We suppose the market to be observable at times $0 = t_0 < t_1 < \cdots < t_N$:

The stock Over each time period $[t_i, t_{i+1}]$ the stock follows the binary model. This is illı in Figure 2.1. After i time periods, the stock can have any of 2^i possible However, *given* its value at time t_i there are only *two* admissible possibilities stock price at time t_{i+1}. It is not necessary, but it is conventional, to suppose time periods have the same length and so we shall write $t_i = i\delta t$ where $\delta t =$

The cash bond In our simple model, the cash bond behaved entirely predictably. There was a interest rate, r, and the cash bond increased in value over a time period of T by a factor e^{rT}. Now, we do not have to impose such a stringent conditic interest rate can itself be random, varying over different time periods. Ou will generalise immediately provided that we insist that the interest rate o time interval $[t_i, t_{i+1})$ is known at the *start* of that interval, although it may on which of the 2^i nodes our market is in. In this way, we admit the pos of randomness in our cash bond. Notice however that it is a very different randomness from that of the stock. The value of the bond at time t_{i+1} is known to us at time t_i. This is certainly not true for the stock. In spite of oı found freedom, for simplicity, we shall continue to suppose that the interest the constant, r.

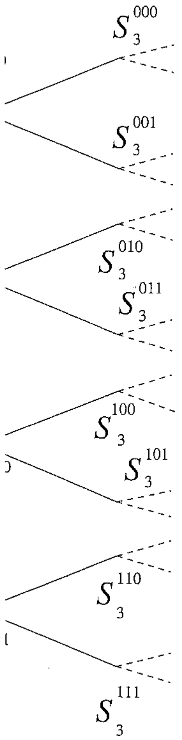

es $0 = t_0 < t_1 < \cdots < t_N = T$.

he binary model. This is illustrated n have any of 2^i possible values. *wo* admissible possibilities for the s conventional, to suppose that all write $t_i = i\delta t$ where $\delta t = T/N$.

ly predictably. There was a known alue over a time period of length se such a stringent condition. The different time periods. Our work sist that the interest rate over the t interval, although it may depend is way, we admit the possibility that it is a very different sort of the bond at time t_{i+1} is already or the stock. In spite of our new- o suppose that the interest rate is

Replicating portfolios

At first sight it is not clear that we can make progress with our new model. For a tree consisting of k time steps there are 2^k possible values for the stock price. If we now look back at Proposition 1.6.5, this suggests that we need at least 2^k stocks to be traded in our market if we want it to be complete. For $k = 20$, this requires over a million 'independent' assets, far more than we see in any real market. But things are not so bad. More claims become attainable if we allow ourselves to rebalance our replicating portfolio after each time period. The only restriction that we impose is that this rebalancing cannot involve any extra input of cash: the purchase of more stock must be funded by the sale of some of our bonds and vice versa. This will be formalised later as the *self-financing* property.

Backwards induction on the tree

The key to understanding pricing and hedging in this bigger model is *backwards induction* on the tree of stock prices.

Example 2.1.1 (Pricing a European call) *Suppose again that we are pricing a European option with maturity time T. As above, we set $\delta t = T/N$ so that T corresponds to N time periods and we write S_i for the stock price at time $i\delta t$. The payoff of the option at time T is denoted by C_N.*

Method: The key idea is as follows. Suppose that we *know* the price, S_{N-1}, of the stock at time $(N-1)\delta t$. Then our previous analysis would tell us the value, C_{N-1}, of the option at time $(N-1)\delta t$. Specifically, $C_{N-1} = \psi_0^{(N)}\mathbb{E}_{N-1}[C_N]$ where the expectation is with respect to a probability measure for which $S_{N-1} = \psi_0^{(N)}\mathbb{E}_{N-1}[S_N]$ and $\psi_0^{(N)} = e^{-r\delta t}$. (In a world of varying interest rates r must be replaced by the rate at the node of the tree corresponding to the *known* value of S_{N-1}.) Moreover, using Lemma 1.3.2, we know how to construct a portfolio at time $(N-1)\delta t$ that will have value exactly C_N at time $N\delta t$. In this way, for each of the 2^{N-1} nodes of the tree at time $(N-1)\delta t$, we calculate the amount of money, C_{N-1}, that we require to construct a portfolio that exactly replicates the claim C_N at time T.

We now think of C_{N-1} as a *claim* at time $(N-1)\delta t$ and we repeat the process. If we *know* S_{N-2}, we can construct a portfolio at time $(N-2)\delta t$ whose value at time $(N-1)\delta t$ will be exactly C_{N-1}, and this portfolio will cost us $\psi_0^{(N-1)}\mathbb{E}_{N-2}[C_{N-1}]$, where the expectation is with respect to a measure such that $S_{N-2} = \psi_0^{(N-1)}\mathbb{E}_{N-2}[S_{N-1}]$. Here again $\psi_0^{(N-1)} = e^{-r\delta t}$. Continuing in this way, we successively calculate the cost of a portfolio that, after appropriate readjustment at each tick of the clock, but without any extra input of wealth and without paying dividends, will allow us to meet exactly the claim against us at time $N\delta t = T$. We'll illustrate the method in Example 2.1.2. □

Binomial trees

It is useful to consider a special form of the binary tree in which over each time step $[t_i, t_{i+1}]$ the stock price either increases from its current value, S_i, to $S_i u$ or decreases to $S_i d$ for some constants $0 < d < u < \infty$. In such a tree the same stock price can be attained in many different ways. For example the value $S_0 ud$ at time t_2 can be attained as the result of an upward stock movement followed by a downward stock

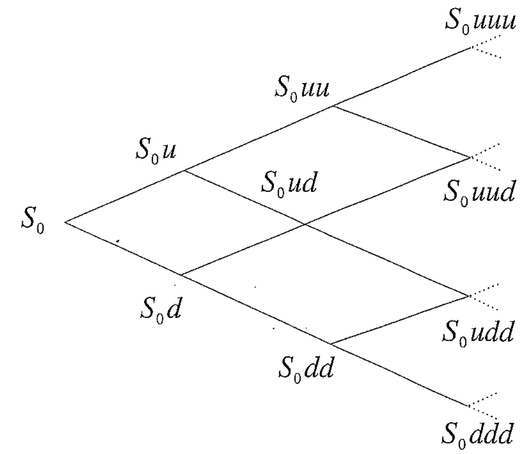

Figure 2.2 A recombinant or binomial tree of stock prices.

movement or vice versa. The tree of stock prices then takes the form of Fi Such a tree is said to be *recombinant* (different branches can recombine special recombinant trees are also known as *binomial trees* since (provided r remain constant over time) the risk-neutral probability measure will be on each upward branch and so the stock price at time $t_n = n\delta t$ is determi binomial distribution. Such trees are computationally much easier to work general binary trees and, as we shall see, are quite adequate for our purpo binomial model was introduced by Cox, Ross & Rubinstein (1979) and has key rôle in the derivatives industry.

We now illustrate the method of backwards induction on a recombinant t

Example 2.1.2 *Suppose that stock prices are given by the tree in Figure that $\delta t = 1$. If interest rates are zero, what is cost of an option to buy the price* 100 *at time* 3*?*

Solution: It is easy to fill in the value of the claim at time 3. Reading fro bottom, the claim has values 60, 20, 0 and 0.

Next we need to find the risk-neutral probabilities for each triad of nod form

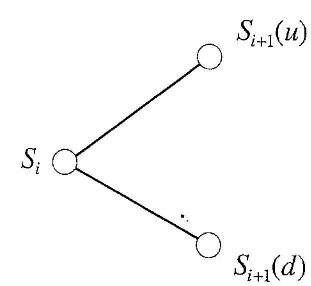

Evidently in this example the risk-neutral probability of stepping up is $1/2$ node. We can now calculate the value of the option at the penultimate time,

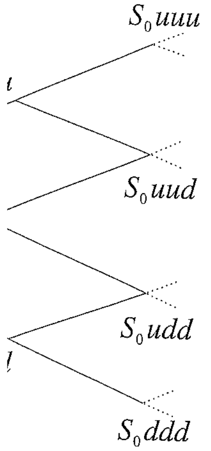

ɜs then takes the form of Figure 2.2. nt branches can recombine). These ɔmial trees since (provided u, d, and ·obability measure will be the same ıt time $t_n = n\delta t$ is determined by a nally much easier to work with than ıite adequate for our purposes. The Rubinstein (1979) and has played a

.duction on a recombinant tree.

given by the tree in Figure 2.3 and ɔst of an option to buy the stock at

im at time 3. Reading from top to

ities for each triad of nodes of the

$S_{i+1}(u)$

$_{i+1}(d)$

lity of stepping up is 1/2 at every ı at the penultimate time, 2, to be,

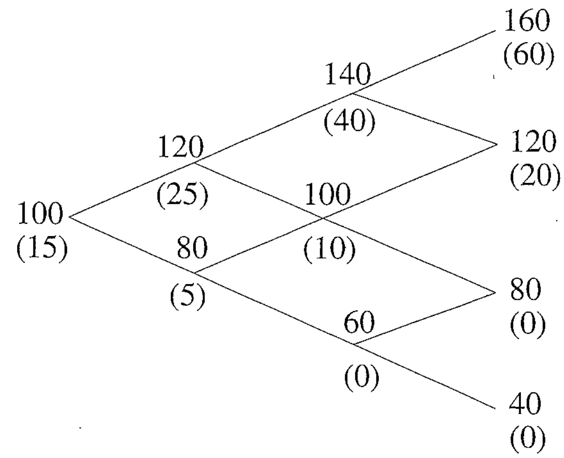

Figure 2.3 The tree of stock prices for the underlying stock in Example 2.1.2. The number in brackets is the value of the claim at each node.

again reading downward, 40, 10, 0. Repeating this for time 1 gives values 25 (if the price steps up from time 0) and 5 (if the price has stepped down). Finally, then, the value of the option at time 0 is 15.

Having filled in the option prices on the tree, we can now construct a portfolio that exactly replicates the claim at time 3 using the prescription of Lemma 1.3.2. We write (ϕ_i, ψ_i) for the amount of stock and bond held in the portfolio over the time interval $[(i-1)\delta t, i\delta t)$.

- At time 0, we are given 15 for the option. We calculate ϕ_1 as $(25-5)/(120-80) = 0.5$. So we buy 0.5 units of stock, which costs 50, and we borrow 35 in cash bonds.
- Suppose that $S_1 = 120$. The new ϕ is $(40-10)/(140-100) = 0.75$, so we buy another 0.25 units of stock, taking our total bond borrowing to 65.
- Suppose that $S_2 = 140$. Now $\phi = (60-20)/(160-120) = 1$, so we buy still more stock, to take our holding up to 1 unit and our total borrowing to 100 bonds.
- Finally, suppose that $S_3 = 120$. The option will be in the money, so we must hand over our unit of stock for 100, which is exactly enough to cancel our bond debt.

The table below summarises our stock and bond holding if the stock price follows another path through the tree.

Time i	Last jump	Stock price S_i	Option value V_i	Stock holding ϕ_i	Bond holding ψ_i
0	—	100	15	—	—
1	down	80	5	0.50	−35
2	up	100	10	0.25	−15
3	down	80	0	0.50	−40

Notice that all of the processes $\{S_i\}_{0\leq i\leq N}$, $\{V_i\}_{0\leq i\leq N}$, $\{\phi_i\}_{1\leq i\leq N}$, $\{\psi_i\}_{1\leq i\leq N}$ on the sequence of up and down jumps. In particular, $\{\phi_i\}_{1\leq i\leq N}$ and $\{\psi_i\}_{1\leq i}$ random too. We do *not* know the dynamics of the portfolio at time 0. However know that our portfolio is *self-financing*. The portfolio that we hold over $[i+1$ can be bought with the proceeds of liquidating (at time $i+1$) the portfolio held over the time interval $[i, i+1)$ – there is no need for any extra input c Moreover, we know how to adjust our portfolio at each time step on the b knowledge of the *current* stock price. There is no risk.

In the single period binary model, we saw that any claim at time T was att and its price at time zero could be expressed as an expectation. The same is the multiperiod setting (see Exercise 1). The proof that any claim is attainable backwards induction on the tree. To recover the pricing formula as an expec we define a probability distribution on paths through the tree.

Path probabilities

Notice that our backwards induction argument has specified exactly one prob on each branch of the tree. For each *path* through the tree that the stock price follow we define the path probability to be the product of the probabilities branches that comprise it.

In Exercise 2 you are asked to show that the price of a claim at time we obtained by backwards induction is precisely the discounted expected va the claim with respect to these path probabilities (in which the discounted c each node is weighted according to the sum of the probabilities of all pat end at that node). Let's just check this prescription for our preceding ex In the recombining tree of Example 2.1.2, there are a total of eight path ending at the top node, one at the bottom and three at each of the other Each path has equal probability, $1/8$, and the expectation of the claim is th $1/8 \times 60 + 3/8 \times 20 = 15$, which is the price that we calculated by bacl induction.

2.2 American options

Our somewhat more sophisticated market model is sufficient for us to take a fir at options whose payoff depends on the *path* followed by the stock price o time interval $[0, T]$. In this section we concentrate on the most important exa of such options: American options.

Definition 2.2.1 (American calls and puts) *An* American call option *with price K and expiry time T gives the holder the right, but not the obligation, an asset for price K at* any time up to T.

An American put option *with strike price K and expiry time T gives the the right, but not the obligation, to* sell *an asset for price K at* any time up to

Evidently the value of an American option should be more than (or at least than) that of its European counterpart. The question is, how much more?

Calls on non-dividend-paying stock

First let us prove the following oft-quoted result.

Lemma 2.2.2 *It is never optimal to exercise an American call option on* non-dividend-*paying stock before expiry.*

Proof: Consider the following two portfolios.

- **Portfolio A:** One American call option plus an amount of cash equal to $Ke^{-r(T-t)}$ at time t.
- **Portfolio B:** One share.

Writing S_t for the share price at time t, if the call option is exercised at time $t < T$, then the value of portfolio **A** at time t is $S_t - K + Ke^{-r(T-t)} < S_t$. (Evidently the option will only be exercised if $S_t > K$.) The value of portfolio **B** is S_t. On the other hand at time T, if the option is exercised then the value of portfolio **A** is $\max\{S_T, K\}$ which is *at least* that of portfolio **B**.

We have shown that exercising prior to maturity gives a portfolio whose value is less than that of portfolio **B** whereas exercising at maturity gives a portfolio whose value is greater than or equal to that of **B**. It cannot be optimal to exercise early. □

This result only holds for non-dividend-paying stock. An alternative proof of Lemma 2.2.2 is Exercise 5. In Exercise 7 the result is extended to show that if the underlying stock pays discrete dividends, then it can only be optimal to exercise at the final time T or at one of the dividend times (see also Exercise 8). More generally, the decision whether to exercise early depends on the 'cost' in terms of lost dividend income.

Put on non-dividend-paying stock

The case of American put options is harder (even without dividends). We illustrate with an example.

Example 2.2.3 *Suppose once again that our asset price evolves according to the recombinant tree of Figure 2.3. To illustrate the method, again we suppose that the risk-free interest rate is zero (but see the second paragraph of Remark 2.2.4). What is the value of a three month American put option with strike price* 100*?*

Solution: As in the case of a European option, we work our way backwards through the tree.

- The value of the claim at time 3, reading from top to bottom, is 0, 0, 20, 60.
- At time 2, we must consider two possibilities: the value if we exercise the claim, and the value if we do not. For the top node it is easy. The value is zero either way. For the second node, the stock price is equal to the strike price, so the value is zero if we exercise the option. On the other hand, if we don't, then from our analysis of the single step binary model, the value of the claim is the expected value under the risk-neutral probabilities of the claim at time 3. We already calculated the risk-neutral

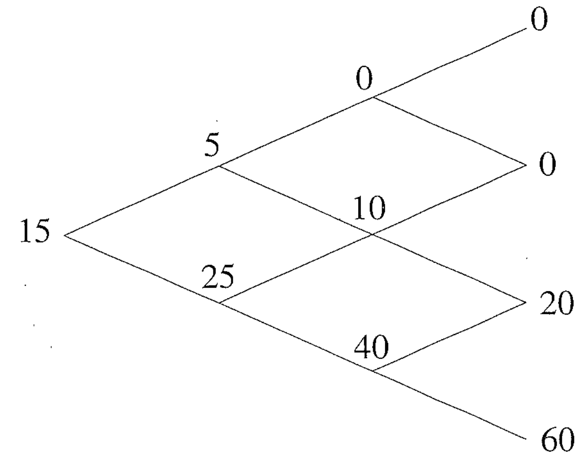

Figure 2.4 The evolution of the price of the American put option of Example 2.2.3.

probabilities to be $1/2$ on each branch of the tree, so this expected value is 10. F the bottom node, the value is 40 whether or not we exercise the claim.

- Now consider the two nodes at time 1. For the top one, if we exercise the option is worthless whereas if we hold it then, again by our analysis of the single peri(model, its value is 5. For the bottom node, if we exercise the option then it is wor 20, whereas if we wait it is worth 25.
- Finally, at time 0, if we exercise, the value is zero, whereas if we wait the value 15.

The option prices are shown in Figure 2.4.

Remark 2.2.4

1. Notice that in the above example it was not optimal to exercise the option at tin 1, even when it was 'in the money'. If $S_1 = 80$, we make 20 from exercisi immediately, but there is 25 to be made from waiting.
2. In this example there was never a strictly positive advantage to early exercise of t option. It was always at least as good to wait. In fact if interest rates are zero this *always* the case, as is shown in Exercise 6. For non-zero interest rates, early exerci can be optimal, see Exercise 9.

2.3 Discrete parameter martingales and Markov processes

Our multiperiod stock market model still looks rather special. To prepare the grou for the continuous time world of later chapters we now place it in the more gene framework of discrete parameter martingales and Markov processes.

First we recall the concepts of random variables and stochastic processes.

Formally, when we talk about a random variable we must first specify a probability triple $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is a set, the *sample space*, $\mathcal{F}$ is a collection of subsets of Ω, *events*, and $\mathbb{P}$ specifies the probability of each event $A \in \mathcal{F}$. The collection $\mathcal{F}$ is a σ-field, that is, $\Omega \in \mathcal{F}$ and $\mathcal{F}$ is closed under the operations of countable union and taking complements. The probability $\mathbb{P}$ must satisfy the usual *axioms of probability:*

- $0 \leq \mathbb{P}[A] \leq 1$, for all $A \in \mathcal{F}$,
- $\mathbb{P}[\Omega] = 1$,
- $\mathbb{P}[A \cup B] = \mathbb{P}[A] + \mathbb{P}[B]$ for any disjoint $A, B \in \mathcal{F}$,
- if $A_n \in \mathcal{F}$ for all $n \in \mathbb{N}$ and $A_1 \subseteq A_2 \subseteq \cdots$ then $\mathbb{P}[A_n] \uparrow \mathbb{P}\left[\bigcup_n A_n\right]$ as $n \uparrow \infty$.

Definition 2.3.1 *A real-valued random variable, X, is a real-valued function on Ω that is $\mathcal{F}$-measurable. In the case of a discrete random variable (that is a random variable that can only take on countably many distinct values) this simply means*

$$\{\omega \in \Omega : X(\omega) = x\} \in \mathcal{F},$$

so that $\mathbb{P}$ assigns a probability to the event $\{X = x\}$. For a general real-valued random variable we require that

$$\{\omega \in \Omega : X(\omega) \leq x\} \in \mathcal{F},$$

so that we can define the distribution function, $F(x) = \mathbb{P}[X \leq x]$.

This looks like an excessively complicated way of talking about a relatively straightforward concept. It is technically required because it may not be *possible* to define $\mathbb{P}$ in a non-trivial way on *all* subsets of Ω, but most of the time we don't go far wrong if we ignore such technical details. However, when we start to study *stochastic processes*, random variables that evolve with time, it becomes much more natural to work in a slightly more formal framework.

To specify a (discrete time) stochastic process, we typically require not just a single σ-field, $\mathcal{F}$, but an increasing sequence of them, $\mathcal{F}_n \subseteq \mathcal{F}_{n+1} \subseteq \cdots \subseteq \mathcal{F}$. The collection $\{\mathcal{F}_n\}_{n\geq 0}$ is then called a *filtration* and the quadruple $\left(\Omega, \mathcal{F}, \{\mathcal{F}_n\}_{n\geq 0}, \mathbb{P}\right)$ is called a *filtered probability space.*

Definition 2.3.2 *A real-valued* stochastic process *is just a sequence of real-valued functions, $\{X_n\}_{n\geq 0}$, on Ω. We say that it is* adapted *to the filtration $\{\mathcal{F}_n\}_{n\geq 0}$ if X_n is $\mathcal{F}_n$-measurable for each n.*

One can then think of the σ-field $\mathcal{F}_n$ as encoding all the information about the evolution of the stochastic process up until time n. That is, if we know whether each event in $\mathcal{F}_n$ happens or not then we can infer the path followed by the stochastic process up until time n. We shall call the filtration that encodes *precisely* this information the *natural* filtration associated to the stochastic process $\{X_n\}_{n\geq 0}$.

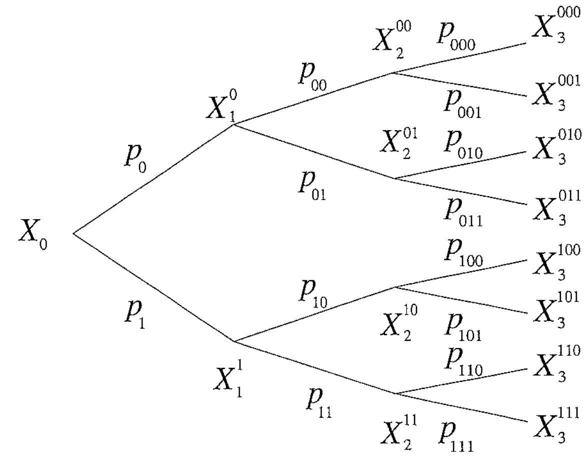

Figure 2.5 Tree representing the stochastic process of Example 2.3.3 and its distribution.

There is an important consequence of the very formal way that this is set u Notice that we have defined the process $\{X_n\}_{n\geq 0}$ as a sequence of measurab functions on Ω *without reference to* $\mathbb{P}$. This is exactly analogous to the situatic in our tree models. We specified the possible values that the stock price could tal at time n, corresponding to prescribing the functions $\{X_n\}_{n\geq 0}$, and superposed tl probabilities afterwards. Even if we had a preconception of what the probabiliti of up and down jumps might be, we then *changed* probability (to the risk-neutr probabilities) in order actually to price claims. This process of changing probabili will be fundamental to our approach to option pricing, even in our most compl market models.

Conditional expectation

When we constructed the probabilities on paths through our binary (or binomia trees, we first specified the probability on each branch of the tree. This was done such a way that the expected value of $e^{-r\delta t}S_{k+1}$ given that the value of the stock time $k\delta t$ is known to be S_k is just S_k. This condition specifies the probabilities on tl two branches emanating from the node corresponding to S_k at time $k\delta t$. We shou like to extend this idea, but first we need to remind ourselves about *condition expectation*. This is best explained through an example.

Example 2.3.3 *Consider the stochastic process represented by the tree in Fi ure 2.5. Its distribution is given by the probabilities on the branches of the tre where, as in §2.1, we assume that the probability of a particular path through t tree is the product of the probabilities of the branches that comprise that path.*

Calculation: Our tree explicitly specifies $\{X_n\}_{n\geq 0}$ and, for a given Ω, implicit specifies $\mathbb{P}$. In later examples we shall be less pedantic, but here we write dov Ω explicitly. There are many possible choices, but an obvious one is the set

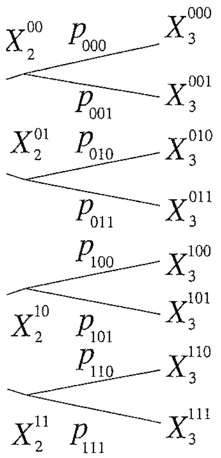

.3.3 and its distribution.

ery formal way that this is set up. $_{n\geq 0}$ as a sequence of measurable exactly analogous to the situation lues that the stock price could take tions $\{X_n\}_{n\geq 0}$, and superposed the nception of what the probabilities ged probability (to the risk-neutral his process of changing probability pricing, even in our most complex

; through our binary (or binomial) ranch of the tree. This was done in given that the value of the stock at on specifies the probabilities on the nding to S_k at time $k\delta t$. We should emind ourselves about *conditional* ample.

ess represented by the tree in Fig- ilities on the branches of the tree, ty of a particular path through the ches that comprise that path.

$_{\geq 0}$ and, for a given Ω, implicitly pedantic, but here we write down but an obvious one is the set of

all possible sequences of 'up' and 'down' jumps. If $\omega = (u, u, d)$, say, then $X_1(\omega) = X_1^0$, $X_2(\omega) = X_2^{00}$ and $X_3(\omega) = X_3^{001}$.

First let us calculate the conditional expectation

$$\mathbb{E}[X_3 | \mathcal{F}_1].$$

Using our interpretation of $\mathcal{F}_n$ as 'information up to time n', our problem is to determine the conditional expectation of X_3 given all the information up to time one. Notice that what we are calculating is an $\mathcal{F}_1$-measurable *random variable*. It depends only on what happened up until time one. There are just two possibilities: the first jump is up, or the first jump is down.

- If the first jump is up, the possible values of X_3 are X_3^{000}, X_3^{001}, X_3^{010} and X_3^{011}. The probability of each value is determined by the path probabilities but restricted to paths emanating from the upper node at time one. The conditional expectation then takes the value

$$\mathbb{E}[X_3 | \mathcal{F}_1](u) = p_{00}p_{000}X_3^{000} + p_{00}p_{001}X_3^{001} + p_{01}p_{010}X_3^{010} + p_{01}p_{011}X_3^{011}.$$

This happens with probability p_0.

- If the first jump is down, which happens with probability p_1, the conditional expectation takes the value

$$\mathbb{E}[X_3 | \mathcal{F}_1](d) = p_{10}p_{100}X_3^{100} + p_{10}p_{101}X_3^{101} + p_{11}p_{110}X_3^{110} + p_{11}p_{111}X_3^{111}.$$

Similarly, we can calculate $\mathbb{E}[X_3 | \mathcal{F}_2]$. This random variable will be $\mathcal{F}_2$-measurable – its value depends on the first two jumps of the process. Its distribution is given in the table below.

Value	Probability
$\mathbb{E}[X\|\mathcal{F}_2](uu) = p_{000}X_3^{000} + p_{001}X_3^{001}$	p_0p_{00}
$\mathbb{E}[X\|\mathcal{F}_2](ud) = p_{010}X_3^{010} + p_{011}X_3^{011}$	p_0p_{01}
$\mathbb{E}[X\|\mathcal{F}_2](du) = p_{100}X_3^{100} + p_{101}X_3^{101}$	p_1p_{10}
$\mathbb{E}[X\|\mathcal{F}_2](dd) = p_{110}X_3^{110} + p_{111}X_3^{111}$	p_1p_{11}

Of course, since $\mathbb{E}[X_3 | \mathcal{F}_2]$ is an $\mathcal{F}_2$-measurable random variable and $\mathcal{F}_1 \subseteq \mathcal{F}_2$, we can calculate the conditional expectation

$$\mathbb{E}[\mathbb{E}[X_3 | \mathcal{F}_2] | \mathcal{F}_1].$$

$$\begin{aligned}
\mathbb{E}[\mathbb{E}[X_3 | \mathcal{F}_2] | \mathcal{F}_1](u) &= p_{00}\mathbb{E}[X_3 | \mathcal{F}_2](uu) + p_{01}\mathbb{E}[X_3 | \mathcal{F}_2](ud), \\
\mathbb{E}[\mathbb{E}[X_3 | \mathcal{F}_2] | \mathcal{F}_1](d) &= p_{10}\mathbb{E}[X_3 | \mathcal{F}_2](du) + p_{11}\mathbb{E}[X_3 | \mathcal{F}_2](dd).
\end{aligned}$$

Substituting the value of $\mathbb{E}[X_3|\mathcal{F}_2]$ from the table above it is easily checked th
this reduces to

$$\mathbb{E}[\mathbb{E}[X_3|\mathcal{F}_2]|\mathcal{F}_1] = \mathbb{E}[X_3|\mathcal{F}_1]. \quad (2.$$

Here is the formal definition.

Definition 2.3.4 (Conditional expectation) *Suppose that X is an $\mathcal{F}$-measurab random variable with $\mathbb{E}[|X|] < \infty$. Suppose that $\mathcal{G} \subseteq \mathcal{F}$ is a σ-field; then t conditional expectation of X given $\mathcal{G}$, written $\mathbb{E}[X|\mathcal{G}]$, is the $\mathcal{G}$-measurable rando variable with the property that for any $A \in \mathcal{G}$*

$$\mathbb{E}[[X|\mathcal{G}]; A] \triangleq \int_A \mathbb{E}[X|\mathcal{G}]\,d\mathbb{P} = \int_A X d\mathbb{P} \triangleq \mathbb{E}[X; A].$$

The conditional expectation exists, but is only unique up to the addition of a rando variable that is zero with probability one. This technical point will be important Exercise 17 of Chapter 3.

Equation (2.1) is a special case of the following key property of condition expectations.

The tower property of conditional expectations: Suppose that $\mathcal{F}_i \subseteq \mathcal{F}_j$; then

$$\mathbb{E}[\mathbb{E}[X|\mathcal{F}_j]|\mathcal{F}_i] = \mathbb{E}[X|\mathcal{F}_i].$$

In words this says that conditioning first on the information up to time j and then c the information up to an earlier time i is the same as conditioning originally up t time i. □

In calculations with conditional expectations, it is often useful to remember th following fact.

Taking out what is known in conditional expectations: Suppose that $\mathbb{E}[X]$ and $\mathbb{E}[XY]$ < ∞; then

$$\text{if } Y \text{ is } \mathcal{F}_n\text{-measurable,} \quad \mathbb{E}[XY|\mathcal{F}_n] = Y\mathbb{E}[X|\mathcal{F}_n].$$

This just says that if Y is known by time n, then if we condition on the informatio up to time n we can treat Y as constant. □

The martingale property

The probability measure on the tree that we used in §2.1 to price claims was chose so that if we define $\{\tilde{S}_k\}_{k\geq 0}$ to be the discounted stock price, that is $\tilde{S}_k = e^{-kr\delta t}S_k$ then the expected value of $\tilde{S}_{k+1}$ given that we know $\tilde{S}_k$ is just $\tilde{S}_k$. We use the notatio

$$\mathbb{E}\big[\tilde{S}_{k+1}|\tilde{S}_k\big] = \tilde{S}_k.$$

Because in our model the stock price has 'no memory', so that the movement of th stock over the next tick of the clock is not influenced by the way in which it reache

its current value, conditioning on knowing $\tilde{S}_k$ is actually the same as conditioning on knowing all of $\mathcal{F}_k$, so that

$$\mathbb{E}\big[\tilde{S}_{k+1}|\mathcal{F}_k\big] = \tilde{S}_k. \tag{2.2}$$

The property (2.2) is sufficiently important that it has a name.

Definition 2.3.5 *Suppose that $\left(\Omega, \{\mathcal{F}_n\}_{n\geq 0}, \mathcal{F}, \mathbb{P}\right)$ is a filtered probability space. The sequence of random variables $\{X_n\}_{n\geq 0}$ is a* martingale *with respect to $\mathbb{P}$ and $\{\mathcal{F}_n\}_{n\geq 0}$ if*

$$\mathbb{E}\left[|X_n|\right] < \infty, \quad \forall n, \tag{2.3}$$

and

$$\mathbb{E}\left[X_{n+1}|\,\mathcal{F}_n\right] = X_n, \quad \forall n. \tag{2.4}$$

If we replace equation (2.4) by

$$\mathbb{E}\left[X_{n+1}|\,\mathcal{F}_n\right] \leq X_n, \quad \forall n,$$

*then $\{X_n\}_{n\geq 0}$ is a $\left(\mathbb{P}, \{\mathcal{F}_n\}_{n\geq 0}\right)$-*super*martingale. If instead we replace it by*

$$\mathbb{E}\left[X_{n+1}|\,\mathcal{F}_n\right] \geq X_n, \quad \forall n,$$

*then $\{X_n\}_{n\geq 0}$ is a $\left(\mathbb{P}, \{\mathcal{F}_n\}_{n\geq 0}\right)$-*sub*martingale.*

These definitions are not exhaustive. There are plenty of processes that fall into none of these categories. A martingale is often thought of as tracking the net gain after successive plays of a fair game. In this setting a supermartingale models net gain from playing an unfavourable game (one we are more likely to lose than to win) and a submartingale is the net gain from playing a favourable game.

It is extremely important to note that the notion of a martingale is really that of a $\left(\mathbb{P}, \{\mathcal{F}_n\}_{n\geq 0}\right)$-martingale. Recall that our definition of stochastic process has divorced the rôles of the sequence $\{\mathcal{F}_n\}_{n\geq 0}$, the $\mathcal{F}$-measurable functions $\{X_n\}_{n\geq 0}$ on Ω and the probability measure $\mathbb{P}$ defined on elements of $\mathcal{F}$. In the setting of §2.1, our view of the market may be that the discounted stock price is *not* a martingale (indeed it probably isn't or no one would ever speculate on stocks – they could get the same money, risk-free, by buying cash bonds). We *change* the probability measure to one which makes the discounted stock price a martingale for the purposes of pricing and, as we shall see, hedging. We shall refer to the probability measure that represents our view of the market as the *market measure*. The new probability measure, which we use for pricing and hedging, is known as the *equivalent martingale measure*.

Remark: (*Martingales indexed by a subset of* $\mathbb{N}$) Although we have defined martingales indexed by $n \in \mathbb{N}$, we shall often talk about martingales indexed by $\{0 \leq n \leq N\}$. They are defined by restricting conditions (2.3) and (2.4) to $\{0 \leq n \leq N\}$. We shall state our key results for martingales indexed by $\{n \geq 0\}$; they can be modified in the obvious way to apply to martingales indexed by $\{0 \leq n \leq N\}$. □

It is often useful to observe that, by the tower property, if $\{X_n\}_{n\geq 0}$ is a $(\mathbb{P}, \{\mathcal{F}_n\}_{n\geq 0}$ martingale then for $i < j$,

$$\mathbb{E}\left[X_j \,|\, \mathcal{F}_i\right] = X_i.$$

The Markov property

Calculations can also be simplified if our martingales have an additional propert the Markov property.

Definition 2.3.6 (Markov process) *The stochastic process $\{X_n\}_{n\geq 0}$ (with its na ural filtration, $\{\mathcal{F}_n\}_{n\geq 0}$) is a discrete time* Markov process *if*

$$\mathbb{P}\left[X_{n+1} \in B \,|\, \mathcal{F}_n\right] = \mathbb{P}\left[X_{n+1} \in B \,|\, X_n\right],$$

for all $B \in \mathcal{F}$.

In words this says that the probability that $X_{n+1} \in B$ given that we know the who history of the process up to time n is the same as the probability that $X_{n+1} \in$ given only the value of X_n. A Markov process has *no memory*. Many of our exampl of martingales (and all our examples of market models) will also have the Markc property. However, not all martingales are Markov processes and not all Markc processes are martingales (see Exercise 11).

> **Notation:** When we wish to emphasise that a filtration is 'generated by' the stochastic process $\{X_n\}_{n\geq 0}$ we use the notation $\{\mathcal{F}_n^X\}_{n\geq 0}$.
> Unless otherwise stated, $\{\mathcal{F}_n\}_{n\geq 0}$ will always be understood to mean the natural filtration associated with the stochastic process under consideration.

It would be excessively pedantic always to insist upon an explicit specification of and so, generally, we won't. We shall also use '$\{X_n\}_{n\geq 0}$ is a $\mathbb{P}$-martingale' to mea $\{X_n\}_{n\geq 0}$ is a $(\mathbb{P}, \{\mathcal{F}_n^X\}_{n\geq 0})$-martingale'.

Examples

Example 2.3.7 (Random walk) *A one-dimensional simple random walk, $\{S_n\}_{n\geq}$ is a Markov process such that $S_{n+1} = S_n + \xi_{n+1}$ where (for each n) $\xi_n \in \{-1, +$ and, under $\mathbb{P}$, $\{\xi_n\}_{n\geq 0}$ are independent identically distributed random variabl Thus*

$$\mathbb{P}\left[S_{n+1} = k+1 \,|\, S_n = k\right] = p, \quad \mathbb{P}\left[S_{n+1} = k-1 \,|\, S_n = k\right] = 1 - p,$$

where $p \in [0, 1]$.

If $p = 0.5$, then $\{S_n\}_{n\geq 0}$ is a $\mathbb{P}$-martingale. If $p < 0.5$ (resp. $p > 0.5$), th $\{S_n\}_{n\geq 0}$ is a $\mathbb{P}$-supermartingale (resp. $\mathbb{P}$-submartingale).

Justification: To check this, notice that since the random walk can be a distance most n from its starting point at time n, the expectation $\mathbb{E}\left[|S_n|\right] < \infty$ is evident

finite. Moreover,

$$\begin{aligned}\mathbb{E}\left[S_{n+1}\mid \mathcal{F}_n\right] &= \mathbb{E}\left[S_n+\xi_{n+1}\mid \mathcal{F}_n\right]\\ &= S_n+\mathbb{E}\left[\xi_{n+1}\mid \mathcal{F}_n\right]\\ &= S_n+\mathbb{E}\left[\xi_{n+1}\right],\end{aligned}$$

where we have used independence of the $\{\xi_n\}_{n\geq 0}$ in the last line. It suffices then to observe that

$$\mathbb{E}\left[\xi_{n+1}\right]\begin{cases} <0, & p<0.5,\\ =0, & p=0.5,\\ >0, & p>0.5.\end{cases}$$

□

Example 2.3.8 (Conditional expectation of a claim) *Suppose that Ω and a filtration $\{\mathcal{F}_n\}_{n\geq 0}$ are given. (The example that we have in mind is that $\mathcal{F}_n$ encodes the history of a financial market up until time $n\delta t$.) Let C_N be any bounded $\mathcal{F}_N$-measurable random variable. (This we are thinking of as a claim against us at time $N\delta t$.) Then for* any *probability measure $\mathbb{P}$, the conditional expectation process, $\{X_n\}_{0\leq n\leq N}$, given by*

$$X_n=\mathbb{E}\left[C_N\mid \mathcal{F}_n\right],$$

is a $\left(\mathbb{P}, \{\mathcal{F}_n\}_{0\leq n\leq N}\right)$-martingale.

Example 2.3.9 (The discounted price of a claim) *In solving our pricing problem for a European option with value C_N at the expiry time $N\delta t$ in the multiperiod binary model of stock prices of §2.1, we found a probability measure, which we denote by $\mathbb{Q}$, under which the discounted stock price is a martingale. For* any *claim, C_N, at time $N\delta t$, provided $\mathbb{E}^{\mathbb{Q}}\left[|C_N|\right]<\infty$, the fair price at time $n\delta t$ of an option with payoff C_N at time $N\delta t$ was found to be*

$$V_n=e^{-r(N-n)\delta t}\mathbb{E}^{\mathbb{Q}}\left[C_N\mid \mathcal{F}_n\right].$$

Define the discounted claim process by $\tilde{V}_n=e^{-rn\delta t}V_n$. Then $\{\tilde{V}_n\}_{0\leq n\leq N}$ is a $\mathbb{Q}$-martingale. This would remain true even if we dropped the assumption of constant interest rates, provided that we knew the risk-free rate over the time interval $[i\delta t, (i+1)\delta t)$ at the beginning *of the period.*

New martingales from old

Our last example shows that the discounted price process of a European option is a martingale. In other words, the discounted value of our replicating portfolio is a martingale. As before, we write (ϕ_n, ψ_n) for the amount of stock and bond held in the replicating portfolio over the nth time interval, that is $[(n-1)\delta t, n\delta t)$. The value of the portfolio at time $n\delta t$ is then

$$V_n=\phi_{n+1}S_n+\psi_{n+1}B_n,$$

where B_n is the value of the cash bond at time $n\delta t$. The portfolio is *self-financing*, that is the cost of constructing the new portfolio at time $(n+1)\delta t$ is exactly offset

by the proceeds of selling the portfolio that we have held over $[n\delta t, (n+1)\delta t)$.] symbols,

$$\phi_{n+1}S_{n+1} + \psi_{n+1}B_{n+1} = \phi_{n+2}S_{n+1} + \psi_{n+2}B_{n+1}.$$

The discounted price is

$$\tilde{V}_n = \phi_{n+1}\tilde{S}_n + \psi_{n+1},$$

and since, using the self-financing property,

$$\phi_{n+1}\tilde{S}_{n+1} + \psi_{n+1} = \phi_{n+2}\tilde{S}_{n+1} + \psi_{n+2},$$

we have

$$\begin{aligned}\tilde{V}_{n+1} - \tilde{V}_n &= \phi_{n+2}\tilde{S}_{n+1} + \psi_{n+2} - \phi_{n+1}\tilde{S}_n - \psi_{n+1}\\ &= \phi_{n+1}\left(\tilde{S}_{n+1} - \tilde{S}_n\right).\end{aligned}$$

That is

$$\tilde{V}_n = V_0 + \sum_{j=0}^{n-1}\phi_{j+1}\left(\tilde{S}_{j+1} - \tilde{S}_j\right). \qquad (2.:$$

From our earlier remarks, $\{\tilde{V}_n\}_{0\leq n\leq N}$ is a $\mathbb{Q}$-martingale, so what we have checke is that under the probability measure $\mathbb{Q}$ for which $\{\tilde{S}_n\}_{0\leq n\leq N}$ is a martingale, tl expression on the right hand side of equation (2.5) is also a martingale. This is pa of a general phenomenon. To state a precise result we need a definition. Recall th we knew ϕ_i at time $(i-1)\delta t$.

Definition 2.3.10 *Given a filtration $\{\mathcal{F}_n\}_{n\geq 0}$, the process $\{A_n\}_{n\geq 1}$ is $\{\mathcal{F}_n\}_{n\geq}$* previsible *or $\{\mathcal{F}_n\}_{n\geq 0}$*-predictable *if A_n is $\mathcal{F}_{n-1}$-measurable for all $n \geq 1$.*

Note that this is the sort of randomness that we have permitted for our cash bond.

Discrete stochastic integrals

Proposition 2.3.11 *Suppose that $\{X_n\}_{n\geq 0}$ is adapted to the filtration $\{\mathcal{F}_n\}_{n\geq 0}$ ar that $\{\phi_n\}_{n\geq 1}$ is $\{\mathcal{F}_n\}_{n\geq 0}$-previsible. Define*

$$Z_n = Z_0 + \sum_{j=0}^{n-1}\phi_{j+1}\left(X_{j+1} - X_j\right), \qquad (2.\text{i}$$

where Z_0 is a constant.

If $\{X_n\}_{n\geq 0}$ is a $\left(\mathbb{P}, \{\mathcal{F}_n\}_{n\geq 0}\right)$-martingale, then so is $\{Z_n\}_{n\geq 0}$.

Remark: If $\{\theta_n\}_{n\geq 0}$ is *adapted* to $\{\mathcal{F}_n\}_{n\geq 0}$, then the process $\{\phi_n\}_{n\geq 1}$ defined by ϕ_n θ_{n-1} is previsible. Thus for an $\{\mathcal{F}_n\}_{n\geq 0}$-adapted process $\{\theta_n\}_{n\geq 0}$, if $\{X_n\}_{n\geq 0}$ is $\left(\mathbb{P}, \{\mathcal{F}_t\}_{t\geq 0}\right)$-martingale then so is

$$Z_n = Z_0 + \sum_{j=0}^{n-1}\theta_j\left(X_{j+1} - X_j\right).$$

have held over $[n\delta t, (n+1)\delta t)$. In

$S_{n+1} + \psi_{n+2} B_{n+1}$.

$_{n+1}$,

$\tilde{S}_{n+1} + \psi_{n+2}$,

$_{+2} - \phi_{n+1}\tilde{S}_n - \psi_{n+1}$

$_n\Big)$.

$_{j+1} - \tilde{S}_j\Big)$. (2.5)

artingale, so what we have checked ich $\{\tilde{S}_n\}_{0\le n\le N}$ is a martingale, the .5) is also a martingale. This is part ult we need a definition. Recall that

, *the process* $\{A_n\}_{n\ge 1}$ *is* $\{\mathcal{F}_n\}_{n\ge 0}$- *measurable for all* $n \ge 1$.

ıave permitted for our cash bond.

ıdapted to the filtration $\{\mathcal{F}_n\}_{n\ge 0}$ *and*

$_{j+1} - X_j\Big)$, (2.6)

ı so is $\{Z_n\}_{n\ge 0}$.

he process $\{\phi_n\}_{n\ge 1}$ defined by $\phi_n =$:d process $\{\theta_n\}_{n\ge 0}$, if $\{X_n\}_{n\ge 0}$ is a

$_{+1} - X_j\Big)$.

□

Proof of Proposition 2.3.11: This is an exercise in the use of conditional expectations.

$$\begin{aligned}
\mathbb{E}\left[Z_{n+1} \middle| \mathcal{F}_n\right] - Z_n &= \mathbb{E}\left[Z_{n+1} - Z_n \middle| \mathcal{F}_n\right] \\
&= \mathbb{E}\left[\phi_{n+1}\left(X_{n+1} - X_n\right) \middle| \mathcal{F}_n\right] \\
&= \phi_{n+1}\mathbb{E}\left[\left(X_{n+1} - X_n\right) \middle| \mathcal{F}_n\right] \\
&= \phi_{n+1}\left(\mathbb{E}\left[X_{n+1} \middle| \mathcal{F}_n\right] - X_n\right) \\
&= 0.
\end{aligned}$$

□

We can think of the sum in equation (2.6) as a *discrete stochastic integral*. When we turn to stochastic integration in Chapter 4, we shall essentially be passing to limits in sums of this form.

The Fundamental Theorem of Asset Pricing

It is not just our binomial models that can be incorporated into the martingale framework. The same argument that allows us to pass from the single period to the multiperiod binary model allows us to pass from the single period models of §1.5 and §1.6 to a multiperiod model. We now recast Theorems 1.5.2 and 1.6.2 in this language. Suppose that our market consists of K stocks and that the possible values that the stock prices $S^1, \dots, S^K$ can take on at times $\delta t, 2\delta t, 3\delta t, \dots, N\delta t = T$ are known. We denote by Ω the set of all possible 'paths' that the stock price vector can follow in $\mathbb{R}_+^K$.

Theorem 1.5.2 tells us that the absence of arbitrage is equivalent to the existence of a probability measure, $\mathbb{Q}$, on Ω that assigns strictly positive mass to every $\omega \in \Omega$ and such that

$$S_{r-1} = \psi_0^{(r)}\mathbb{E}^{\mathbb{Q}}[S_r | S_{r-1}],$$

where S_r is the vector of stock prices at time r and $\psi_0^{(r)}$ is the discount on riskless borrowing over $[(r-1)\delta t, r\delta t]$.

If, as above, we consider the *discounted* stock prices, $\{\tilde{S}_j\}_{0\le j\le N}$, given by $\tilde{S}_j = \prod_{i=1}^{j}\psi_0^{(i)} S_j$, then

$$\mathbb{E}^{\mathbb{Q}}[\tilde{S}_r | \tilde{S}_1, \dots, \tilde{S}_{r-1}] = \mathbb{E}^{\mathbb{Q}}\left[\tilde{S}_r \middle| \mathcal{F}_{r-1}\right] = \tilde{S}_{r-1}.$$

In other words, the discounted stock price vector is a $\mathbb{Q}$-*martingale.*

Definition 2.3.12 *Two probability measures* $\mathbb{P}$ *and* $\mathbb{Q}$ *on a space* Ω *are said to be* equivalent *if for all events* $A \subseteq \Omega$

$$\mathbb{Q}(A) = 0 \qquad \textit{if and only if} \qquad \mathbb{P}(A) = 0.$$

Suppose then that we have a market model in which the stock price vector can follow one of a finite number of paths Ω through $\mathbb{R}_+^K$. We may even have our own belief as to how the price will evolve, encoded in a probability measure, $\mathbb{P}$, on Ω. Theorem 1.5.2 and Theorem 1.6.2 combine to say:

Theorem 2.3.13 *For the multiperiod market model described above, there is arbitrage if and only if there is an equivalent martingale measure $\mathbb{Q}$. That is, th is a measure, $\mathbb{Q}$, equivalent to $\mathbb{P}$, such that the discounted stock price process i $\mathbb{Q}$-martingale.*

In that case, the time zero market price of an attainable claim C_N (to be delive at time $N\delta t$) is unique and is given by

$$\mathbb{E}^{\mathbb{Q}}[\psi_0 C_N],$$

where $\psi_o = \prod_1^N \psi_0^{(i)}$ is the discount factor over N periods.

Although there are extra technical conditions, this fundamental theorem has ess tially the same statement for markets that evolve continuously with time.

2.4 Some important martingale theorems

Phrasing everything in the martingale framework places many powerful theore at our disposal. In this section, we present some of the most important result the theory of discrete parameter martingales. However, our coverage is necessa cursory. An excellent and highly readable account is Williams (1991).

Stopping times

One of the most important calculational tools in martingale theory is the Optic Stopping Theorem. Before we can state it, we need to introduce the notion c stopping time.

Definition 2.4.1 *Given a sample space Ω equipped with a filtration $\{\mathcal{F}_n\}_{n\geq}$* stopping time *or* optional time *is a random variable $T : \Omega \to \mathbb{Z}_+$ with the prop that*

$$\{T \leq n\} \in \mathcal{F}_n, \qquad \textit{for all } n \geq 0.$$

This just says that we can decide whether or not $T \leq n$ on the basis of the informa available at time n – we don't need to look into the future.

Example 2.4.2 *Consider the simple random walk of Example 2.3.7. Define T be the first time that the random walk takes the value 1, that is*

$$T = \inf\{i \geq 0 : S_i = 1\};$$

then T is a stopping time.

On the other hand,

$$U = \sup\{i \geq 0 : S_i = 1\}$$

is not *a stopping time.*

Optional stopping

An equivalent definition of stopping time is that the random variable $\theta_n \triangleq \mathbf{1}_{\{T\geq n}$ for $n \geq 0$, is *adapted* (see Definition 2.3.2). Consequently, from the remark follow

Proposition 2.3.11, if $\{X_n\}_{n\geq 0}$ is a martingale, then so is the process

$$Z_n \triangleq \sum_{j=0}^{n-1} \theta_j \left(X_{j+1} - X_j\right). \tag{2.7}$$

Notice that we can rearrange this expression,

$$\begin{aligned} Z_n &= \sum_{j=0}^{n-1} \theta_j \left(X_{j+1} - X_j\right) \\ &= \sum_{j=0}^{n-1} \mathbf{1}_{\{T\geq j+1\}} \left(X_{j+1} - X_j\right) \\ &= X_{T\wedge n} - X_0, \end{aligned}$$

where $T \wedge n$ denotes the minimum of T and n.

Theorem 2.4.3 (Optional Stopping Theorem) *Let $\left(\Omega, \mathcal{F}, \{\mathcal{F}_n\}_{n\geq 0}, \mathbb{P}\right)$ be a filtered probability space. Suppose that the process $\{X_n\}_{n\geq 0}$ is a $\left(\mathbb{P}, \{\mathcal{F}_n\}_{n\geq 0}\right)$-martingale, and that T is a* bounded *stopping time. Then*

$$\mathbb{E}\left[\, X_T \mid \mathcal{F}_0\right] = X_0,$$

and hence

$$\mathbb{E}\left[X_T\right] = X_0.$$

Proof: The proof is a simple application of the calculation that we did above. If we know that $T \leq N$, then in the notation of (2.7), $Z_N = X_T - X_0$ and since $\{Z_n\}_{n\geq 0}$ is a martingale, $\mathbb{E}\left[\, Z_N \mid \mathcal{F}_0\right] = Z_0 = 0$, i.e.

$$\mathbb{E}\left[\, X_T \mid \mathcal{F}_0\right] = X_0.$$

Taking expectations once again yields

$$\mathbb{E}\left[X_T\right] = X_0.$$

$\square$

It is essential in this result that the stopping time be *bounded*. In practice this will be the case in all of our financial applications, but Exercise 15 shows what can go wrong. More general versions of the theorem are available; see for example, Williams (1991). Here we satisfy ourselves with an application (see also Exercise 14).

Proposition 2.4.4 *Let $\{S_n\}_{n\geq 0}$ be the (asymmetric) simple random walk of Example 2.3.7 with $p > 1/2$. For $x \in \mathbb{Z}$ we write*

$$T_x = \inf\{n : S_n = x\},$$

and define

$$\phi(x) = \left(\frac{1-p}{p}\right)^x.$$

Then for $a < 0 < b$,

$$\mathbb{P}[T_a < T_b] = \frac{1-\phi(b)}{\phi(a)-\phi(b)}.$$

Proof: We first show that $\{\phi(S_n)\}_{n\geq 0}$ is a $\mathbb{P}$-martingale. Since the walk can only tal one step at a time, $-n \leq S_n \leq n$. Using also that $0 < (1-p)/p < 1$ for $p > 1/$ we evidently have that

$$\mathbb{E}[|\phi(S_n)|] < \infty, \qquad \forall n.$$

To check that we really have a martingale is reduced to another exercise conditional expectations. We must calculate

$$\mathbb{E}\left[\phi(S_{n+1}) | \mathcal{F}_n\right].$$

Recall that $S_{n+1} = \sum_{j=1}^{n+1} \xi_j = S_n + \xi_{n+1}$, where, under $\mathbb{P}$, the random variables are independent and identically distributed with

$$\mathbb{P}[\xi_j = 1] = p \quad \text{and} \quad \mathbb{P}[\xi_j = -1] = 1-p.$$

This gives

$$\begin{aligned}
\mathbb{E}\left[\phi(S_{n+1}) | \mathcal{F}_n\right] &= \mathbb{E}\left[\phi(S_n)\left(\frac{1-p}{p}\right)^{\xi_{n+1}}\middle| \mathcal{F}_n\right] \\
&= \phi(S_n)\mathbb{E}\left[\left(\frac{1-p}{p}\right)^{\xi_{n+1}}\right] \\
&= \phi(S_n)\left(p\left(\frac{1-p}{p}\right)^1 + (1-p)\left(\frac{1-p}{p}\right)^{-1}\right) \\
&= \phi(S_n).
\end{aligned}$$

We should now like to apply the Optional Stopping Theorem to the stopping tir $T = T_a \wedge T_b$, the first time that the walk hits either a or b. The difficulty is that T not *bounded*. Instead then, we apply the theorem to the stopping time $T \wedge N$ for arbitrary (deterministic) N. This gives

$$\begin{aligned}
1 &= \mathbb{E}[\phi(S_0)] = \mathbb{E}[\phi(S_{T\wedge N})] \\
&= \phi(a)\mathbb{P}[S_T = a, T \leq N] + \phi(b)\mathbb{P}[S_T = b, T \leq N] + \mathbb{E}[\phi(S_N), T > N
\end{aligned} \quad (2.$$

Now

$$\begin{aligned}
0 \leq \mathbb{E}[\phi(S_N), T > N] &= \mathbb{E}[\phi(S_N) | T > N]\,\mathbb{P}[T > N] \\
&\leq \left[\left(\frac{1-p}{p}\right)^b + \left(\frac{p}{1-p}\right)^a\right]\mathbb{P}[T > N],
\end{aligned}$$

and since $\mathbb{P}[T > N] \to 0$ as $N \to \infty$, we can let $N \to \infty$ in (2.8) to deduce that

$$\phi(a)\mathbb{P}[S_T = a] + \phi(b)\mathbb{P}[S_T = b] = 1. \tag{2.9}$$

Finally, since $\mathbb{P}[S_T = a] = 1 - \mathbb{P}[S_T = b]$, and $\mathbb{P}[T_a < T_b] = \mathbb{P}[S_T = a]$, equation (2.9) becomes

$$\phi(a)\mathbb{P}[T_a < T_b] + \phi(b)\left(1 - \mathbb{P}[T_a < T_b]\right) = 1.$$

Rearranging,

$$\mathbb{P}[T_a < T_b] = \frac{1 - \phi(b)}{\phi(a) - \phi(b)},$$

as required. □

A convergence theorem

Often one can deduce a great deal about martingales from apparently scant information. An example is the result of Exercise 12 which says that a previsible martingale is constant. Another example is provided by the following result.

Theorem 2.4.5 (Positive Supermartingale Convergence Theorem) *If $\{X_n\}_{n\geq 0}$ is a $(\mathbb{P}, \{\mathcal{F}_n\}_{n\geq 0})$-supermartingale and $X_n \geq 0$ for all n, then there exists an $\mathcal{F}_\infty$-measurable random variable, X_∞, with $\mathbb{E}[X_\infty] < \infty$ such that with $\mathbb{P}$-probability one*

$$X_n \to X_\infty \quad \textit{as } n \to \infty.$$

A proof of this result is beyond our scope here, but can be found, for example, in Williams (1991).

Compensation

Before returning to some finance, we record just one more result. Recall that submartingales tend to rise on the average and supermartingales fall on the average. The following result, sometimes called *compensation*, says that we can subtract a non-decreasing process from a submartingale to obtain a martingale and we can add a non-decreasing process to a supermartingale to obtain a martingale. In both cases, the interesting thing is that the non-decreasing processes are *previsible*.

Proposition 2.4.6

1. *Suppose that $\{X_n\}_{n\geq 0}$ is a $(\mathbb{P}, \{\mathcal{F}_n\}_{n\geq 0})$-submartingale. Then there is a previsible, non-decreasing process $\{A_n\}_{n\geq 0}$ such that $\{X_n - A_n\}_{n\geq 0}$ is a $(\mathbb{P}, \{\mathcal{F}_n\}_{n\geq 0})$-martingale. If we insist that $A_0 = 0$, then $\{A_n\}_{n\geq 0}$ is unique.*
2. *Suppose that $\{X_n\}_{n\geq 0}$ is a $(\mathbb{P}, \{\mathcal{F}_n\}_{n\geq 0})$-supermartingale. Then there is a previsible, non-decreasing process $\{A_n\}_{n\geq 1}$ such that $\{X_n + A_n\}_{n\geq 0}$ is a $(\mathbb{P}, \{\mathcal{F}_n\}_{n\geq 0})$-martingale. If we insist that $A_0 = 0$, then $\{A_n\}_{n\geq 0}$ is unique.*

Proof: The proofs of the two parts are essentially identical, so we restrict our attention to 1.

Define $A_0 = 0$ and then for $n \geq 1$ set

$$A_n - A_{n-1} = \mathbb{E}\left[X_n - X_{n-1} \middle| \mathcal{F}_{n-1}\right].$$

By definition $\{A_n\}_{n\geq 0}$ will be previsible and non-decreasing (since $\{X_n\}_{n\geq 0}$ i submartingale). We must check that $\{X_n - A_n\}_{n\geq 0}$ is a martingale. First we che that $\mathbb{E}[|X_n - A_n|] < \infty$ for all n.

$$\begin{aligned}
\mathbb{E}[|X_n - A_n|] &\leq \mathbb{E}[|X_n|] + \mathbb{E}[A_n] \\
&= \mathbb{E}[|X_n|] + \mathbb{E}\left[A_0 + \sum_{j=1}^{n}(A_j - A_{j-1})\right] \\
&= \mathbb{E}[|X_n|] + \sum_{j=1}^{n}\mathbb{E}\left[\mathbb{E}\left[X_j - X_{j-1} \middle| \mathcal{F}_{j-1}\right]\right] \quad \text{(by definition of } A \\
&\leq \mathbb{E}[|X_n|] + \sum_{j=1}^{n}\mathbb{E}\left[\mathbb{E}\left[|X_j| + |X_{j-1}| \middle| \mathcal{F}_{j-1}\right]\right] \\
&= \mathbb{E}[|X_n|] + \sum_{j=1}^{n}\mathbb{E}\left[|X_j| + |X_{j-1}|\right] \quad \text{(tower property)},
\end{aligned}$$

and evidently this final expression is finite since by assumption $\mathbb{E}\left[|X_j|\right] < \infty$ f all j.

Now we check the martingale property,

$$\begin{aligned}
&\mathbb{E}\left[X_{n+1} - A_{n+1} \middle| \mathcal{F}_n\right] \\
&\quad = \mathbb{E}\left[X_{n+1} - \mathbb{E}\left[X_{n+1} - X_n \middle| \mathcal{F}_n\right] - A_n \middle| \mathcal{F}_n\right] \quad \text{(by definition of } A_{n+1}\text{)} \\
&\quad = \mathbb{E}\left[X_{n+1} - X_{n+1} + X_n - A_n \middle| \mathcal{F}_n\right] \\
&\quad = X_n - A_n.
\end{aligned}$$

It remains to check that if $A_0 = 0$ then the process $\{A_n\}_{n\geq 0}$ is unique. Suppose th there were another predictable process $\{B_n\}_{n\geq 0}$ with the same property. Then $\{X_n$ $A_n\}_{n\geq 0}$ and $\{X_n - B_n\}_{n\geq 0}$ are both martingales and, therefore, so is the differen between them, $\{A_n - B_n\}_{n\geq 0}$. On the other hand $\{A_n - B_n\}_{n\geq 0}$ is predictable ar *predictable martingales are constant* (see Exercise 12). Since $A_0 = 0 = B_0$, tl proof is complete.

American options and supermartingales

Let's see what these concepts correspond to in a financial example.

Example 2.4.7 (American options revisited) *Assume the binomial model and n tation of §2.2 and let $\mathbb{Q}$ be the probability measure on the tree under which tl discounted stock price $\{\tilde{S}_n\}_{0\leq n\leq N}$ is a martingale. We denote by $\{\tilde{V}_n\}_{0\leq n\leq N}$ tl discounted value of an American call or put option with strike K and maturi $T = N\delta t$ and define*

$$\tilde{B}_n = \begin{cases} e^{-n\delta t}(S_n - K)_+ & \text{in the case of the call,} \\ e^{-n\delta t}(K - S_n)_+ & \text{in the case of the put.} \end{cases}$$

$_{n-1} | \mathcal{F}_{n-1}]$.

on-decreasing (since $\{X_n\}_{n\geq 0}$ is a $_{\geq 0}$ is a martingale. First we check

$- A_{j-1})\Big]$

$_{j-1} | \mathcal{F}_{j-1}]]$ (by definition of A_j)

$|X_{j-1}| | \mathcal{F}_{j-1}]]$

$_{-1}|]$ (tower property),

by assumption $\mathbb{E}\left[|X_j|\right] < \infty$ for

$| \mathcal{F}_n]$ (by definition of A_{n+1})

ss $\{A_n\}_{n\geq 0}$ is unique. Suppose that vith the same property. Then $\{X_n -$ and, therefore, so is the difference d $\{A_n - B_n\}_{n\geq 0}$ is predictable and ise 12). Since $A_0 = 0 = B_0$, the □

financial example.

ssume the binomial model and no- asure on the tree under which the ale. We denote by $\{\tilde{V}_n\}_{0\leq n\leq N}$ the ption with strike K and maturity

e case of the call,
e case of the put.

(The filtration is always that generated by $\{S_n\}_{0\leq n\leq N}$.) Then $\{\tilde{V}_n\}_{0\leq n\leq N}$ is the smallest $\mathbb{Q}$-supermartingale that dominates $\{\tilde{B}_n\}_{0\leq n\leq N}$.

In Exercise 16 it is shown that this characterisation provides yet another simple proof of Lemma 2.2.2.

Explanation for example: We know from §2.2 that

$$\tilde{V}_{n-1} = \max\left\{\tilde{B}_{n-1}, \mathbb{E}^{\mathbb{Q}}\left[\tilde{V}_n \middle| \mathcal{F}_{n-1}\right]\right\}, \quad 0 \leq n \leq N,$$

and $\tilde{V}_N = \tilde{B}_N$. Evidently $\{\tilde{V}_n\}_{0\leq n\leq N}$ is a supermartingale that dominates $\{\tilde{B}_n\}_{0\leq n\leq N}$. To check that it is the *smallest* supermartingale with this property, suppose that $\{\tilde{U}_n\}_{0\leq n\leq N}$ is any other supermartingale that dominates $\{\tilde{B}_n\}_{n\geq 0}$. Then $\tilde{U}_N \geq \tilde{V}_N$, and if $\tilde{U}_n \geq \tilde{V}_n$, then

$$\tilde{U}_{n-1} \geq \mathbb{E}^{\mathbb{Q}}\left[\tilde{U}_n \middle| \mathcal{F}_{n-1}\right] \geq \mathbb{E}^{\mathbb{Q}}\left[\tilde{V}_n \middle| \mathcal{F}_{n-1}\right],$$

and so

$$\tilde{U}_{n-1} \geq \max\left\{\tilde{B}_{n-1}, \mathbb{E}^{\mathbb{Q}}\left[\tilde{V}_n \middle| \mathcal{F}_{n-1}\right]\right\} = \tilde{V}_{n-1}.$$

The result follows by backwards induction. The process $\{\tilde{V}_n\}_{0\leq n\leq N}$ is called the *Snell envelope* of $\{\tilde{B}_n\}_{0\leq n\leq N}$. □

Remark: Proposition 2.4.6 tells us that we can write

$$\tilde{V}_n = \tilde{M}_n - \tilde{A}_n$$

where $\{\tilde{M}_n\}_{n\geq 0}$ is a martingale and $\{\tilde{A}_n\}_{n\geq 0}$ is a non-decreasing process, with $A_0 = 0$. Since the market is complete, we can hedge M_N exactly by holding a portfolio that consists over the nth time step of ϕ_n units of stock and ψ_n units of cash bond. The seller of the American option would more than meet her liability by holding such a portfolio. The holder of the option will exercise at the first time j when $\tilde{A}_{j+1}$ is non-zero (recall that the process $\{\tilde{A}_n\}_{n\geq 0}$ is *previsible*), since at that time it is better to sell the option and invest the money according to the hedging portfolio $\{(\phi_n, \psi_n)\}_{j\leq n\leq N}$. □

2.5 The Binomial Representation Theorem

Pricing a derivative in the martingale framework corresponds to taking an expectation. But arbitrage prices are only meaningful if we can construct a hedging portfolio. If we know the hedging portfolio then we saw in the discussion preceding Definition 2.3.10 that we can express the discounted value of the portfolio, and therefore of the derivative, as a 'discrete stochastic integral' of the stock holding in the portfolio with respect to the discounted stock price. In order to pass from the discounted price of the derivative to a hedging portfolio we need the following converse to Proposition 2.3.11. We work in the context of our binomial model of stock prices.

Theorem 2.5.1 (Binomial Representation Theorem) *Suppose that the measur is such that the discounted binomial price process* $\{\tilde{S}_n\}_{n\geq 0}$ *is a* $\mathbb{Q}$*-marting If* $\{\tilde{V}_n\}_{n\geq 0}$ *is any other* $(\mathbb{Q}, \{\mathcal{F}_n\}_{n\geq 0})$*-martingale, then there exists an* $\{\mathcal{F}_n\}_r$ *predictable process* $\{\phi_n\}_{n\geq 1}$ *such that*

$$\tilde{V}_n = \tilde{V}_0 + \sum_{j=0}^{n-1} \phi_{j+1}\left(\tilde{S}_{j+1} - \tilde{S}_j\right). \qquad (2.$$

Proof: We consider a single time step for our binomial tree. It is convenient to w

$$\Delta\tilde{V}_{i+1} = \tilde{V}_{i+1} - \tilde{V}_i \quad \text{and} \quad \Delta\tilde{S}_{i+1} = \tilde{S}_{i+1} - \tilde{S}_i.$$

Given their values at time $i\delta t$, each of $\tilde{V}_{i+1}$ and $\tilde{S}_{i+1}$ can take on one of two poss values that we denote by $\{\tilde{V}_{i+1}(u), \tilde{V}_{i+1}(d)\}$ and $\{\tilde{S}_{i+1}(u), \tilde{S}_{i+1}(d)\}$ respectively

We should like to write $\Delta\tilde{V}_{i+1} = \phi_{i+1}\Delta\tilde{S}_{i+1} + k_{i+1}$, where ϕ_{i+1} and k_{i+1} both known at time $i\delta t$. In other words we seek ϕ_{i+1} and k_{i+1} such that

$$\tilde{V}_{i+1}(u) - \tilde{V}_i = \phi_{i+1}\left(\tilde{S}_{i+1}(u) - \tilde{S}_i\right) + k_{i+1},$$

and

$$\tilde{V}_{i+1}(d) - \tilde{V}_i = \phi_{i+1}\left(\tilde{S}_{i+1}(d) - \tilde{S}_i\right) + k_{i+1}.$$

Solving this gives

$$\phi_{i+1} = \frac{\tilde{V}_{i+1}(u) - \tilde{V}_{i+1}(d)}{\tilde{S}_{i+1}(u) - \tilde{S}_{i+1}(d)}$$

and $k_{i+1} = \tilde{V}_{i+1}(u) - \tilde{V}_i - \phi_{i+1}\left(\tilde{S}_{i+1}(u) - \tilde{S}_i\right)$, both of which are known at ti $i\delta t$.

Now $\{\tilde{V}_i\}_{i\geq 0}$ and $\{\tilde{S}_i\}_{i\geq 0}$ are both martingales so that

$$\mathbb{E}\left[\Delta\tilde{V}_{i+1}\middle|\mathcal{F}_i\right] = 0 = \mathbb{E}\left[\Delta\tilde{S}_{i+1}\middle|\mathcal{F}_i\right]$$

from which it follows that $k_{i+1} = 0$.

In other words,

$$\Delta\tilde{V}_{i+1} = \phi_{i+1}\Delta\tilde{S}_{i+1},$$

where ϕ_{i+1} is known at time $i\delta t$. Induction ties together all these increments into result that we want.

From martingale representation to replicating portfolio

From our previous work, we know that if $\{\tilde{V}_i\}_{i\geq 0}$ is the discounted price of a cla then such a predictable process $\{\phi_i\}_{i\geq 1}$ arises as the stock holding when we constr our replicating portfolio. We should like to go the other way. Given $\{\phi_i\}_{i\geq 1}$, can construct a self-financing replicating portfolio? Not surprisingly, the answer is ye

) *Suppose that the measure* $\mathbb{Q}$
rocess $\{\tilde{S}_n\}_{n\geq 0}$ *is a* $\mathbb{Q}$*-martingale.*
gale, then there exists an $\{\mathcal{F}_n\}_{n\geq 0}$-

$$\tilde{S}_{j+1} - \tilde{S}_j\Big). \qquad (2.10)$$

inomial tree. It is convenient to write

$$\Delta\tilde{S}_{i+1} = \tilde{S}_{i+1} - \tilde{S}_i.$$

$\tilde{S}_{i+1}$ can take on one of two possible
d $\{\tilde{S}_{i+1}(u), \tilde{S}_{i+1}(d)\}$ respectively.
$_{+1} + k_{i+1}$, where ϕ_{i+1} and k_{i+1} are
ϕ_{i+1} and k_{i+1} such that

$$u) - \tilde{S}_i\Big) + k_{i+1},$$

$$d) - \tilde{S}_i\Big) + k_{i+1}.$$

$$\frac{\tilde{V}_{i+1}(d)}{\tilde{S}_{i+1}(d)}$$

$_i\Big)$, both of which are known at time

s so that

$$\left[\Delta\tilde{S}_{i+1}\middle|\mathcal{F}_i\right]$$

$$\tilde{S}_{i+1},$$

ogether all these increments into the

□

$_0$ is the discounted price of a claim,
the stock holding when we construct
ne other way. Given $\{\phi_i\}_{i\geq 1}$, can we
Not surprisingly, the answer is yes.

Construction strategy: At time i, buy a portfolio that consists of ϕ_{i+1} units of stock and $\tilde{V}_i - \phi_{i+1}\tilde{S}_i$ units of cash bond.

We must check that this strategy really works. It is convenient to write B_i for the value of the bond at time $i\delta t$.

Suppose that at time $i\delta t$ we have bought ϕ_{i+1} units of stock and $\left(\tilde{V}_i - \phi_{i+1}\frac{S_i}{B_i}\right)$ units of cash bond. This will cost us

$$\phi_{i+1}S_i + \left(\tilde{V}_i - \phi_{i+1}\frac{S_i}{B_i}\right)B_i = \tilde{V}_i B_i = V_i.$$

The value of this portfolio at time $(i+1)\delta t$ is then

$$\begin{aligned}\phi_{i+1}S_{i+1} + \left(\tilde{V}_i - \phi_{i+1}\frac{S_i}{B_i}\right)B_{i+1} &= B_{i+1}\left(\phi_{i+1}\left(\frac{S_{i+1}}{B_{i+1}} - \frac{S_i}{B_i}\right) + \tilde{V}_i\right)\\ &= \tilde{V}_{i+1}B_{i+1} \quad \text{(by the binomial representation)}\\ &= V_{i+1},\end{aligned}$$

which is exactly enough to construct our new portfolio at time $(i+1)\delta t$. Moreover, at time $N\delta t$ we have precisely the right amount of money to meet the claim against us.

> **Three steps to replication:** There are three steps to pricing and hedging a claim C_T against us at time T.
>
> - Find a probability measure $\mathbb{Q}$ under which the discounted stock price (with its natural filtration) is a martingale.
> - Form the discounted value process,
>
> $$\tilde{V}_i = e^{-ri\delta t}V_i = \mathbb{E}^{\mathbb{Q}}\left[e^{-rT}C_T\middle|\mathcal{F}_i\right].$$
>
> - Find a predictable process $\{\phi_i\}_{1\leq i\leq N}$ such that
>
> $$\Delta\tilde{V}_i = \phi_i\Delta\tilde{S}_i.$$

2.6 Overture to continuous models

Before rigorously deriving the acclaimed Black–Scholes pricing formula for the value of a European option, we are going to develop a substantial body of material. As an appetiser though, we can use our discrete techniques to see what form our results must take in the continuous world.

It is easy to believe that we should be able to use a discrete model with very small time periods to approximate a continuous model. The Black–Scholes model is based on the lognormal model that we mentioned in §1.2. With this in mind, we choose our approximation to have constant growth rate and constant 'noise'.

Model with constant stock growth and noise

The model is parametrised by the time period, δt, and three fixed constant parar ters, ν, σ and the riskless rate r.

• The cash bond has the form $B_t = e^{rt}$, which does not depend on the inter size.

• The stock price process follows the nodes of a binomial tree. If the current vá of the stock is s, then over the next time period it moves to the new value

$$\begin{cases} s\exp\left(\nu\delta t + \sigma\sqrt{\delta t}\right) & \text{if up,} \\ s\exp\left(\nu\delta t - \sigma\sqrt{\delta t}\right) & \text{if down.} \end{cases}$$

Suppose our belief is that the jumps are equally likely to be up or down. So un the *market* measure, $\mathbb{P}\left[\text{up jump}\right] = 1/2 = \mathbb{P}\left[\text{down jump}\right]$ at each time step.

For a fixed time t, set N to be the number of time periods until time t, tha $N = t/\delta t$. Then

$$S_t = S_0 \exp\left(\nu t + \sigma\sqrt{t}\left(\frac{2X_N - N}{\sqrt{N}}\right)\right),$$

where X_N is the total number of the N separate jumps which were up jumps. To what happens as $\delta t \to 0$ (or equivalently $N \to \infty$) we call on the Central Li Theorem.

Theorem 2.6.1 (Central Limit Theorem) *Let $\xi_1, \xi_2, \ldots$ be a sequence of indep dent identically distributed random variables under the probability measure $\mathbb{P}$ v finite mean μ and finite non-zero variance σ^2 and let $S_n = \xi_1 + \ldots + \xi_n$. Then*

$$\frac{S_n - n\mu}{\sqrt{n\sigma^2}}$$

converges in distribution to an $N(0, 1)$ random variable as $n \to \infty$.

Now X_N is the sum of N independent random variables $\{\xi_i\}_{1 \le i \le N}$ taking the vε $+1$ with probability $\frac{1}{2}$ and 0 otherwise. This means $\mathbb{E}\left[\xi_i\right] = \frac{1}{2}$ and $var\left[\xi_i\right] = \frac{1}{4}$ that by the Central Limit Theorem, the distribution of the random variable $(2X_N$ $N)/\sqrt{N}$ converges to that of a normal random variable with mean zero and varia one. In other words, as δt gets smaller (and so N gets larger), the distribution o converges to that of a lognormal distribution. More precisely, in the limit, $\log S$ normally distributed with mean $\log S_0 + \nu t$ and variance $\sigma^2 t$.

Under the martingale measure

This is what happens under the original measure $\mathbb{P}$. What happens under martingale measure, $\mathbb{Q}$, that we use for pricing?

By Lemma 1.3.2, under the martingale measure, the probability of an up jum

$$p = \frac{\exp(r\delta t) - \exp(\nu\delta t - \sigma\sqrt{\delta t})}{\exp(\nu\delta t + \sigma\sqrt{\delta t}) - \exp(\nu\delta t - \sigma\sqrt{\delta t})},$$

which is approximately

$$\frac{1}{2}\left(1 - \sqrt{\delta t}\left(\frac{\nu + \frac{1}{2}\sigma^2 - r}{\sigma}\right)\right).$$

So under the martingale measure, $\mathbb{Q}$, X_N is still binomially distributed, but now has mean Np and variance $Np(1-p)$.

Thus, under $\mathbb{Q}$, $(2X_N - N)/\sqrt{N}$ has mean that tends to $-\sqrt{t}\left(\nu + \frac{1}{2}\sigma^2 - r\right)/\sigma$ and variance that approaches one as δt tends to zero. Again using the Central Limit Theorem the random variable $(2X_N - N)/\sqrt{N}$ converges to a normally distributed random variable, with mean $-\sqrt{t}\left(\nu + \frac{1}{2}\sigma^2 - r\right)/\sigma$ and variance one. Under $\mathbb{Q}$ then, S_t is lognormally distributed with mean $\log S_0 + (r - \frac{1}{2}\sigma^2)t$ and variance $\sigma^2 t$. This can be written

$$S_t = \exp\left(\sigma\sqrt{t}Z + \left(r - \frac{1}{2}\sigma^2\right)t\right),$$

where, under $\mathbb{Q}$, the random variable Z is normally distributed with mean zero and variance one.

Pricing a call option

If our discrete theory carries over to the continuous limit, then in our continuous model the price at time zero of a European call option with strike price K at time T will be the discounted expected value of the claim under the martingale measure, that is

$$\mathbb{E}^{\mathbb{Q}}\left[e^{-rT}\left(S_T - K\right)_+\right],$$

where r is the riskless rate. Substituting, we obtain

$$\mathbb{E}^{\mathbb{Q}}\left[\left(S_0 \exp\left(\sigma\sqrt{T}Z - \frac{1}{2}\sigma^2 T\right) - K\exp\left(-rT\right)\right)_+\right]. \tag{2.11}$$

We'll derive this pricing formula rigorously in Chapter 5 where we'll also show that equation (2.11) can be evaluated as

$$S_0\Phi\left(\frac{\log\frac{S_0}{K} + \left(r + \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}\right) - Ke^{-rT}\Phi\left(\frac{\log\frac{S_0}{K} + \left(r - \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}\right),$$

where Φ is the standard normal distribution function,

$$\Phi(z) = \mathbb{Q}\left[Z \le z\right] = \int_{-\infty}^{z} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx.$$

Exercises

1 Notice that, like the single period ternary model of Chapter 1, the two-step binomial model allows the stock to take on three distinct values at time 2. Show, however, that every claim can be exactly replicated by a self-financing portfolio, that is, the market is *complete*.

More generally, show that if the market evolves according to a k-step binomial model then it is complete.

2 Show that the price of a claim obtained by backwards induction on the binomial t is precisely the value obtained by calculating the discounted expected value of claim with respect to the path probabilities introduced in §2.1.

3 Consider two dates T_0, T_1 with $T_0 < T_1$. A *forward start option* is a contrac which the holder receives at time T_0, at no extra cost, an option with expiry date and strike price equal to S_{T_0} (the asset price at time T_0). Assume that the stock pı evolves according to a two-period binary model, in which the asset price at time is either S_0u or S_0d, and at time T_1 is one of S_0u^2, S_0ud and S_0d^2 with

$$d < \min\left\{e^{rT_0}, e^{r(T_1-T_0)}\right\} \leq \max\left\{e^{rT_0}, e^{r(T_1-T_0)}\right\} < u,$$

where r denotes the risk-free interest rate. Find the fair price of such an optioı time zero.

4 A *digital option* is one in which the payoff depends in a discontinuous way on asset price. The simplest example is the *cash-or-nothing option*, in which the pay to the holder at maturity T is $X\mathbf{1}_{\{S_T > K\}}$ where X is some prespecified cash sum. Suppose that an asset price evolves according to the binomial model in which each step, the asset price moves from its current value S_n to one of S_nu and S_nd. usual, if ΔT denotes the length of each time step, $d < e^{r\Delta T} < u$.
Find the time zero price of the above option. You may leave your answer as a sun

5 Let C_t denote the value at time t of an American call option on non-dividend-pay stock with strike price K and maturity T. If the risk-free interest rate is $r > 0$, prc that

$$C_t \geq S_t - Ke^{-r(T-t)} > S_t - K,$$

and deduce that it is never optimal to exercise this option prior to the maturity tir T.

6 Let C_t be as in Exercise 5 and let P_t be the value of an American put option on t same stock with the same strike price and maturity. By comparing the values of t suitable portfolios, show that

$$C_t + K \geq P_t + S_t.$$

Using put–call parity for European options and the result of Exercise 5, show that

$$P_t \geq C_t + Ke^{-r(T-t)} - S_t.$$

Combine these results to see that, if $r > 0$ and $t < T$,

$$S_t - K \leq C_t - P_t < S_t - Ke^{-r(T-t)}$$

and deduce that if interest rates are zero, there is no advantage to early exercise the put.

kwards induction on the binomial tree ; the discounted expected value of the roduced in §2.1.

forward start option is a contract in tra cost, an option with expiry date T_1 time T_0). Assume that the stock price lel, in which the asset price at time T_0 ou^2, S_0ud and S_0d^2 with

$$\mathrm{x}\left\{e^{rT_0}, e^{r(T_1-T_0)}\right\} < u,$$

nd the fair price of such an option at

epends in a discontinuous way on the *r-nothing option*, in which the payoff X is some prespecified cash sum. to the binomial model in which, at it value S_n to one of S_nu and S_nd. As ep, $d < e^{r\Delta T} < u$. ou may leave your answer as a sum.

n call option on non-dividend-paying risk-free interest rate is $r > 0$, prove

$$> S_t - K,$$

this option prior to the maturity time,

lue of an American put option on the rity. By comparing the values of two

$$+ S_t.$$

the result of Exercise 5, show that

$$^{-t)} - S_t.$$

$$< T,$$

$$- Ke^{-r(T-t)}$$

is no advantage to early exercise of

7 If a stock price is S just before a dividend D is paid, what is its value immediately after the payment? Suppose that a stock pays dividends at discrete times, $T_0, T_1, \dots, T_n$. Show that it can be optimal to exercise an American call on such a stock *prior* to expiry.

8 Suppose that the stock in Figure 2.3 will pay a dividend of 5% of its value at time 2. As before, interest rates are zero and between times 2 and 3 the value of the stock will either increase or decrease by 20. Find the time zero price of an American call option on this stock with strike 100 and maturity 3. Is it ever optimal to exercise early?

9 Consider the American put option of Example 2.2.3, but now suppose that interest rates are such that a \$1 cash bond at time $i\delta t$ is worth \$1.1 at time $(i+1)\delta t$. Find the value of the put. At what time will it be exercised?

10 Suppose that an asset price evolves according to the binomial model. For simplicity suppose that the risk-free interest rate is zero and ΔT is 1. Suppose that under the probability $\mathbb{P}$, at each time step, stock prices go up with probability p and down with probability $1-p$.
The conditional expectation

$$M_n \triangleq \mathbb{E}[S_N|\mathcal{F}_n], \qquad 1 \leq n \leq N,$$

is a stochastic process. Check that it is a $\mathbb{P}$-martingale and find the distribution of the random variable M_n.

11 (a) Find a Markov process that is not a martingale.
(b) Find a martingale that is not a Markov process.

12 Show that a previsible martingale is constant.

13 Let $\{S_n\}_{n\geq 0}$ be simple random walk under the measure $\mathbb{P}$. Calculate $\mathbb{E}[S_n]$ and $var[S_n]$.

14 Let $\{S_n\}_{n\geq 0}$ be a symmetric simple random walk under the measure $\mathbb{P}$, that is, in the notation of Example 2.3.7, $p = 1/2$. Show that $\{S_n^2\}_{n\geq 0}$ is a $\mathbb{P}$-submartingale and that $\{S_n^2 - n\}_{n\geq 0}$ is a $\mathbb{P}$-martingale.
Let $T = \inf\{n : S_n \notin (-a, a)\}$, where $a \in \mathbb{N}$. Use the Optional Stopping Theorem (applied to a suitable sequence of bounded stopping times) to show that $\mathbb{E}[T] = a^2$.

15 As in Exercise 14, let $\{S_n\}_{n\geq 0}$ be a symmetric simple random walk under $\mathbb{P}$ and write $X_n = S_n + 1$. (Note that $\{X_n\}_{n\geq 0}$ is a simple random walk started from 1 at time zero.)
Let $T = \inf\{n : X_n = 0\}$. Show that T is a stopping time and that if $Y_n = X_{T\wedge n}$, then $\{Y_n\}_{n\geq 0}$ is a non-negative martingale and therefore, by Theorem 2.4.5, converges to a limit, Y_∞ as $n \to \infty$.
Show that $\mathbb{E}[Y_n] = 1$ for all n, but that $Y_\infty = 0$. Why does this not contradict the conclusion of the Optional Stopping Theorem?

16 Recall Jensen's inequality: if g is a convex function and X a real-valued ranc variable then

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

Combine this with the characterisation (Example 2.4.7) of the discounted p of an American call option on non-dividend-paying stock as the smallest supermartingale that dominates $\{e^{-rn\delta t}(S_n - K)_+\}_{n\geq 0}$ to prove that the price o American call on non-dividend-paying stock is the same as that of a European with the same strike and maturity.