Chapter 3

Binomial asset pricing models

It is not uncommon for analysts to model stock prices in discrete time. Usually these models are simpler than those which are continuous in time. Additionally, the time scale of discrete models can be made to match the data collected from the real world. For instance, in Australia's recently reformed electricity market spot prices for electricity only exist at five minute intervals, not continuously. Another example is the majority of private stock market investors, those who only have access to daily closing prices.

Binomial models are particularly simple because they assume that at each step stock price movements are limited to only two possible values. These models are especially attractive because of the ease at which they are implemented numerically. Furthermore, the discrete-time results obtained in the binomial tree framework often correspond to the continuous-time results in the limit as the time step tends to zero.

From a learning perspective, binomial asset pricing models provide a good setting to play with martingales and risk-neutral pricing.

3.1 A single step model

Suppose the market consists of two securities: bonds which may be used for riskless borrowing and lending, and a single type of stock. Unlimited amounts of either can be bought and sold without incurring transaction costs. The stock has an initial price

of s_0 and one time tick, δt, later will be worth $S_{\delta t}$. $S_{\delta t}$ is a random variable which takes one of two possible values; s_u if the stock moves up, or s_d if the stock moves down. Probabilities may be assigned to these stock movements. In summary

$$S_{\delta t} = \begin{cases} s_u, & \text{with probability } p; \\ s_d, & \text{w.p. } 1 - p. \end{cases}$$

Here's an example which fits nicely into this framework.

Example 4 (Roulette). *One possible type of bet at the roulette table in a casino is to place a \$1 chip on an odd number coming up. This event has a probability of $\frac{18}{37}$ and pays \$2. Before the wheel is spun another punter asks if you'll sell a call option on such a bet with strike price $K = 50c$ for $C_0 = 73c$. This option would be worth $2 - 0.50$ if an odd number appears, and nothing in the event of any other number. A quick calculation indicates the expected value of C_T is*

$$\mathbb{E}(C_T) = (2 - 0.5) \times \frac{18}{37} + 0 \times \frac{19}{37} \approx \$0.7297.$$

So 73c might be a fair price, but is it really what the contract is worth?
I claim that at this price it is always possible to make a risk-free profit.

The risk-free profit may be realised using the following strategy. Sell an odd number bet worth 75c, buy the call costing 73c, and keep 2c in hand. After the wheel is spun there are only two possible outcomes:

odd number *need to pay \$1.50 to cover the bet, but the call is now worth $2 - 0.5$ dollars - exactly cancelling this out. We still have 2c in hand.*

any other number *both the bet and the call are worth zero, and we still have a 2c profit.*

So this transaction will result in risk-free profit for the buyer of the call option. The price of the contract is set too low, resulting in an arbitrage opportunity.

This example illustrates that the call option is not priced at its expectation. Even if the stock moves with transition probabilities which are known, pricing using these probabilities will, in general, be wrong.

The correct price is found using a result mentioned in Section 2.1. *The unique initial price is equal to the cost of setting up a portfolio which will replicate the claim.* In other words, if we can construct a portfolio (in terms of stocks and cash bonds) which has the same payoff function as the contract in question, then the cost of the contract at time zero will be equal to the initial cost to set up the portfolio. If it weren't then there would be an arbitrage opportunity. One could buy (or sell as appropriate) the replicating portfolio against the option, wait for maturity and simply collect the difference.

As an example we shall construct the price of the call option. Recall that (in our artificial universe) \$1 now is worth \$$e^{rT}$, T time units later. Also recall that the payoff of a call option at maturity is

$$C_T = \begin{cases} S_T - K, & \text{if } S_T > K; \\ 0, & \text{otherwise.} \end{cases}$$

Consider a portfolio consisting of ϕ units of stock and ψ worth of cash bond. It costs $V_0 = \phi s_0 + \psi$ to set up this portfolio at time zero. One time-step later, though, there are two cases to consider:

$$V_T = \begin{cases} \phi s_u + \psi e^{rT}, & \text{if step up;} \\ \phi s_d + \psi e^{rT}, & \text{if step down.} \end{cases}$$

We are aiming to replicate C_T, so set

$$s_u - K = \phi s_u + \psi e^{rT},$$
$$\text{and}$$
$$0 = \phi s_d + \psi e^{rT},$$

to form two linear equations in two unknowns; ϕ and ψ. The solution is given by

$$\phi = \frac{s_u - K}{s_u - s_d},$$
$$\psi = e^{-rT} \left(\frac{-s_d(s_u - K)}{s_u - s_d} \right).$$

If we bought this (ϕ, ψ) portfolio at time zero, the construction ensures that it will be worth the same as the call option at time T regardless of step direction. Therefore, the value of the call option at time zero must be the initial cost of the portfolio. Namely

$$V_0 = s_0 \left(\frac{s_u - K}{s_u - s_d} \right) + e^{-rT} \left(\frac{-s_d(s_u - K)}{s_u - s_d} \right).$$

We have derived the single step price for a call option. If it were any other price there would be an arbitrage opportunity to make money. A price for which no arbitrage opportunity exists is said to be *risk-neutral*.

Let's apply this to the Roulette example.

Example 5 (Roulette revisited). *Recall $s_0 = 1$ and S_T was either 0 or 2. The interest is negligible since the payment occurs immediately. The strike price K for the call option was 0.50 and we are interested in determining the amount the contract is worth initially, C_0. Using the solution above we have*

$$\phi = \frac{2 - 0.50}{2 - 0},$$
$$\psi = \frac{-0(2 - 0.50)}{2 - 0},$$

and the contract must be worth 75c.

3.2 A multiple step model

As a model for a real stock the single step binary branch has obvious limitations. It is important as a building block for our next section on multi-period binomial models; in other words, binomial trees. The simplicity of these models allows the results to be obtained explicitly, yet rich enough to approximate many realistic situations.

Still working only with cash bonds and a single stock we shall construct a hedging strategy on a binomial tree. Crucial to our arguments will be the assumption of no transaction fees as this allows us to readjust our portfolio at each time tick.

Here is a summary of the binomial tree model.

two securities a risky asset (stock) and a riskless asset (cash bond);

time a series of times $0, \delta t, 2\,\delta t, \ldots, N\,\delta t = T$ at which trades can occur;

interest the risk-free interest rate fluctuates between time steps with the interest rate for the next period known at the start of that period. Let r_i be the risk-free interest rate for the period $[(i-1)\delta t, i\delta t)$. Under this assumption \$1 now grows to be worth $B_T = \prod_{i=1}^{N} e^{r_i \delta t}$ at time T. Similarly, \$1 promised at some future time T is worth $B_T^{-1} = e^{-\sum_{i=1}^{N} r_i \delta t}$ now;

binomial tree a binomial tree of possible states for the stock prices. As in the single period model, suppose that at each time tick, the stock moves from its current value, along a branch, to one of two other values. The stock is assumed to have an initial value S_0 which is known and constant. After n time ticks the value of the stock price S_n could be any one of 2^n possibilities depending on the random choice of path taken.

The famous Cox-Ross-Rubinstein model (Merton 1992, Musiela & Rutkowski 1997) is a binomial tree with a special property. At each step of the CRR model the stock price may move from s_{now} to either $s_u = u\,s_{\text{now}}$ or $s_d = d\,s_{\text{now}}$, where d and u are constants which satisfy $0 < d < 1 + r < u$.

It is common to assign a set of transition probabilities, $\mathbb{P}$. Under $\mathbb{P}$ each branch is attached a weight corresponding to the tendency of the price process to travel that branch. The roulette example of Section 3.1 demonstrated that stock derivatives are not priced using these probabilities. It turns out that if there exists an arbitrage price for a claim then there must be a measure $\mathbb{Q}$, separate from $\mathbb{P}$, which also assigns probabilities to each of the transitions, and the value of the claim *is* the expectation under $\mathbb{Q}$.

The following example is from Baxter & Rennie (1996):

Example 6. *Suppose the price process of a particular stock wanders on a binary tree. Figure 3.1 gives the possible values for the first three steps. For simplicity we assume*
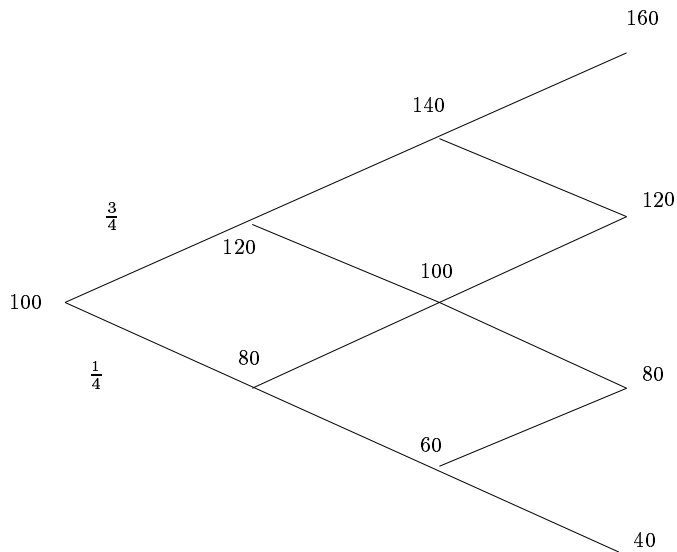
Figure 3.1: Stock price

the interest rate is zero. The transition probabilities are known to be $\frac{3}{4}$ for a step up and $\frac{1}{4}$ for down. For this example the probability that the stock after the third step is worth \$120 is

$$\Pr(S_3 = 120) = \frac{3}{4}\frac{3}{4}\frac{1}{4} + \frac{3}{4}\frac{1}{4}\frac{3}{4} + \frac{1}{4}\frac{3}{4}\frac{3}{4},$$

since there are three paths which lead to 120. The expected value of S_3 under this measure is $\mathbb{E}_{\mathbb{P}}(S_3) = 130$ (check this). What is the value of a European call option to buy stock for \$100 at time 3?

Recalling that the value of a call is $\max(S_T - K, 0)$ at time of maturity, we can start filling in an analogous tree for the value movements of a call. The four time 3 nodes will be simply an application of the payoff function: $\max(S_3 - K, 0)$.

For all times earlier than 3 we apply the construction strategy of Section 3.1; at time 2 suppose the current stock price is 140, and that we hold ϕ units of stock and ψ cash bond. Note that in the next time tick the stock price will move to either 160 or 120, and the contract will be worth 60 or 20. We want our portfolio to match the
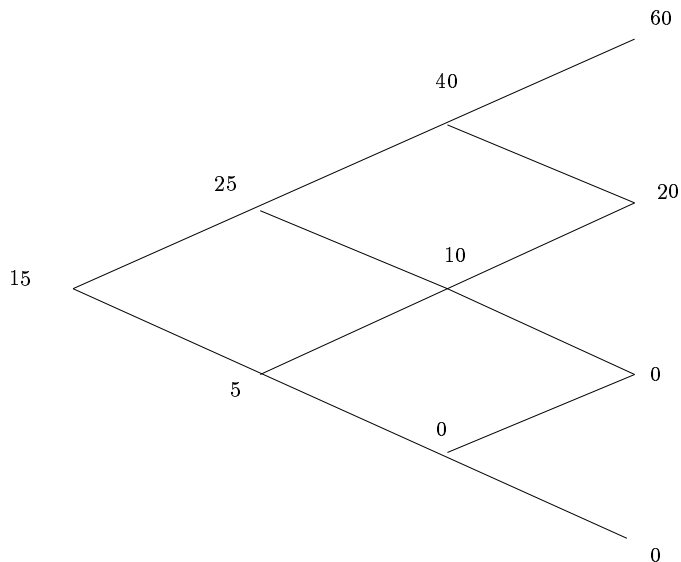
Figure 3.2: Value of the call option

value of the call, so set

$$60 = 160\phi + \psi, \tag{3.1}$$
$$20 = 120\phi + \psi.$$

Solving this set of linear equations yields

$$\phi = \frac{60 - 20}{160 - 120},$$

$$\psi = 60 - 160\left(\frac{60 - 20}{160 - 120}\right).$$

The value of the call at $t = 2$ if $S_2 = 140$ must be

$$C_2 = 140\phi + \psi = 40. \tag{3.2}$$

Figure 3.2 gives the value of the call option at each step on the binary tree.

Solving equations analogous to (3.1) and plugging the solution into (3.2) for each node is cumbersome. Some careful rearranging (see question 3) reveals these equations

can be replaced by

$$V_{\text{now}} = e^{-r\delta t}\left(qV_{\text{up}} + (1-q)V_{\text{down}}\right), \tag{3.3}$$

where

$$q = \frac{e^{r\delta t}s_{\text{now}} - s_d}{s_u - s_d},$$

and where V_{now} denotes the current contract's value, and V_{up} (resp. V_{down}) is future value of the contract if the stock moves up to s_u (resp. down to s_d). This formula holds for the put and other options as well.

Notice that no reference to the transition probabilities $\mathbb{P} = (p, 1 - p)$ was made in deriving the solution. Instead we have $\mathbb{Q} = (q, 1 - q)$ which assigns also weights to each branch. Furthermore, using Equation (3.3) it is easy to see that the initial premium paid for the option is the discounted expected payoff taken with respect to $\mathbb{Q}$,

$$V_0 = e^{-rT}\mathbb{E}_{\mathbb{Q}}(V_T). \tag{3.4}$$

The values, ϕ and ψ, which were obtained at each step are still useful. They specify a trading strategy which will perfectly hedge the contract. The numbers ϕ and ψ are the quantities of stocks and bonds which should be held for the next time step to replicate the value of the contract given the current stock price. They can be obtained at each tick time using the solution to the equations analogous to (3.1) for the general case; namely

$$V_{\text{up}} = s_u\phi + \psi,$$
$$V_{\text{down}} = s_d\phi + \psi.$$

The next example uses the stock price process and associated contract prices described in Example 6 (and pictured in Figures 3.1 and 3.2) to describe the trading strategy a merchant bank might use. The bank sells a call option to an investor at time zero, their goal now is to design a strategy (in terms of bonds and stocks) which will successfully replicate the value of the contract. By doing this the bank avoids any exposure to risk.

Example 7. *The replicating portfolio starts with*

$$\phi = (25 - 5)/(120 - 80) = 0.50,$$
$$\psi = 15 - 0.5 \times 100 = -35.$$

The original portfolio should contain 0.5 *units of stock and* -35 *dollars. The bank sells the option for* 15 *dollars and should borrow another* 35 *and use this* 50 *dollars to buy half a unit of stock worth* 0.5 × 100.

Stock moves down to 80. *The new replicating portfolio contains*

$$\phi = (10 - 0)/(100 - 60) = 0.25,$$
$$\psi = 5 - 0.25 \times 80 = -15.$$

Sell a one quarter share of stock worth 0.25 × 80 = 20 *dollars. This reduces our debt to* 15 *dollars.*

Stock moves up to 100. *The new portfolio contains*

$$\phi = (20 - 0)/(120 - 80) = 0.50,$$
$$\psi = 10 - 0.5 \times 100 = -40.$$

We need to buy a quarter share of stock, this costs 0.25 × 100 = 25 *dollars. Our debt is now* 15 + 25 = 40 *dollars.*

Stock moves up to 120. *Our holdings now include a* 0.50 *share of stock and a* 40 *dollar debt. The call finished in the money so we need to by another* 0.50 *units of stock for* 0.5 × 120 = 60 *dollars, and sell all the stock at the strike price* 100 *dollars. This* 100 *exactly cancels our debt* 40 + 60. *The strategy results in neither loss nor a gain.*

Notice that the portfolio changes from one time to the next but the changes are *self-financing*. That is to say, the total value of the portfolio before and after each

trade are the same. The bank neither receives nor spends money except when it initially sells the option.

The example shows the importance of tracking ϕ, the number of units of stock to be held as you leave a given node. It characterises the replicating (hedging) portfolio.

3.3 Risk-neutral pricing

In this section we shall investigate Formula (3.4),

$$V_0 = e^{-rT}\mathbb{E}_{\mathbb{Q}}(V_T),$$

in a little more depth. This relationship between $\mathbb{Q}$, the random payoff V_T, and the contract's initial value V_0 is not an attribute of the binomial tree model. Instead it is a consequence of the arbitrage-free market environment our model assumes and the risk-neutral price which we are seeking.

The measures $\mathbb{Q}$ and $\mathbb{P}$ are nothing more than collections of weights assigned to each branch of the binomial tree. The subjective measure $\mathbb{P}$ describes how likely each transition is from each node. It is chosen according to assumptions about the process being modelled. On the other hand, $\mathbb{Q}$ is imposed by the underlying stock price process, and has nothing to do with any distributional assumptions. For the binomial tree we found $\mathbb{Q}$ by construction. At each step q was the weight of a step up and $1 - q$ for a step down. We found

$$q = \frac{e^{r\delta t}s_{\text{now}} - s_d}{s_u - s_d}.$$

What is special about $\mathbb{Q}$ to make Equation (3.4) hold, and does it carry over to other models? If the belated background on conditional expectation, filtrations, and martingales, can be excused, the answer is given by the following theorem.

Theorem 3. *A market is arbitrage-free iff there exists a unique martingale[1] measure,* $\mathbb{Q}$ *for the discounted stock price process* $B_t^{-1}S_t$*. In this case, the risk-neutral price at*

---

[1]See the appendix to this chapter for background on conditional expectation, filtrations, and martingales.

time t of any European style claim with payoff V_T is given by the risk-neutral valuation formula

$$V_t(\mathcal{F}_t) = B_t \mathbb{E}_\mathbb{Q}(B_T^{-1} V_T | \mathcal{F}_t). \tag{3.5}$$

The implications of this result are significant. For one, it provides an easy method to check whether a given market model is arbitrage-free: if there is exactly one probability measure under which the discounted stock price process $B_t^{-1} S_t$ is a martingale, then we know the market is arbitrage-free. More importantly, when Equation (3.5) can be evaluated, it provides a closed-form expression for the time t value of any contingent claim. As enticing as this sounds it is only useful if the model is simple enough that the right hand side of Equation (3.5) can be evaluated. Much of the time spent modelling financial markets is concentrated on searching for a happy medium between realism and tractability.

Example 8 (The contingent claim which pays S_T). *If the payoff of the contract is equal to the value of the stock at expiry, the claim can be replicated by simply purchasing one unit of stock - therefore the value of the contract now must be S_0. Let's verify that this is the price we would obtain using Theorem 3.*

We are aiming to show $V_t(\mathcal{F}_t) = S_t$ at each time-tick t. Without loss of generality, take the time between steps, δt, to be 1. Then V_n, S_n, and $\mathcal{F}_n$ the values of the contract, the underlying stock, and the process history at time $n \, \delta t$. Similarly $B_t = B_n$ when $t = n \, \delta t$.

Consider the contract value a single time-step before expiry. At time $T - \delta t = N - 1$,

$$V_{N-1}(\mathcal{F}_{N-1}) \;\; = \;\; B_{N-1} \, \mathbb{E}_\mathbb{Q}(B_N^{-1} \, V_N \, | \, \mathcal{F}_{N-1}),$$

$$= \;\; e^{-r_N}[q s_u + (1 - q)s_d]. \tag{3.6}$$

Substituting

$$q = \frac{e^{r_N} s_{N-1} - s_d}{s_u - s_d}$$

into Equation (3.6) *confirms* $V_{N-1}(\mathcal{F}_{N-1}) = s_{N-1}$.

   Proceeding backwards by induction, at any other time-tick $n\,\delta t$

$$V_n(\mathcal{F}_n) \;=\; B_n\,\mathbb{E}_{\mathbb{Q}}(B_N^{-1}V_N \,|\, \mathcal{F}_n),$$

$$\begin{aligned}
\;=\; & B_n \big[q\,\mathbb{E}_{\mathbb{Q}}\left(B_N^{-1}V_N \,|\, \mathcal{F}_n, S_{n+1} = s_u\right) \\
& + (1-q)\,\mathbb{E}_{\mathbb{Q}}\left(B_N^{-1}V_N \,|\, \mathcal{F}_n, S_{n+1} = s_d\right)\big],
\end{aligned} \tag{3.7}$$

where

$$q = \frac{e^{r_{n+1}}s_n - s_d}{s_u - s_d}.$$

Now, under the inductive assumption

$$V_{n+1}(\mathcal{F}_{n+1}) = B_{n+1}\,\mathbb{E}_{\mathbb{Q}}\left(B_N^{-1}V_N \,|\, \mathcal{F}_{n+1}\right) = S_{n+1},$$

and so (3.7) *becomes*

$$V_n(\mathcal{F}_n) = \frac{B_n}{B_{n+1}}[qs_u + (1-q)s_d] = S_n.$$

Example 9 (Forward contract with strike price K). *Consider a forward contract maturing at $T = N\,\delta t$. This contract costs nothing to enter into and has payoff $V_T = S_T - K$. Using Theorem 3 we know $V_0 = \mathbb{E}_{\mathbb{Q}}(B_T^{-1}V_T)$, which can be evaluated in terms of K, s_0, and B_T. If the initial contract cost V_0 is 0, then the strike price must be fixed at*

$$K = e^{\sum_{k=1}^{N} r_k\,\delta t}\, s_0.$$

   The discrete-time models we've considered so far have required some fairly unrealistic assumptions. In particular, if the stock price process wanders on a binomial tree, then price changes only occur at certain times those changes are restricted to be in some finite set of possibilities.

   In the next section we will look at a model which takes the set of future stock prices to be $\mathbb{R}^+$; any positive real number. The model is called the *log-normal model of stock price dynamics* and may be approximated using binomial trees. The Black-Scholes formulae for pricing European call and put options are heuristically derived in the discrete time limit $\delta t \to 0$.

3.4 Log-normal price dynamics

At the time when a contract is written we don't know the stock price at maturity, we can only guess at it. The models of the preceding section assumed future stock prices are part of some finite set with path probabilities representing the likelihood of achieving each price. Another widely accepted model is that stock prices are log-normally distributed. That is, there are constants μ and σ such that the log of the stock price at time T divided by that at time 0, $\log(S_T/S_0)$, is normally distributed with mean μT and variance $\sigma^2 T$. In symbols:

$$
\begin{aligned}
\Pr\left(\frac{S_T}{S_0} \in [a, b]\right) &= \Pr\left(\log\left(\frac{S_T}{S_0}\right) \in [\log a, \log b]\right), \\
&= \int_{\log a}^{\log b} \frac{1}{\sqrt{2\pi T}\sigma} \exp\left(-\frac{(x - \mu T)^2}{2\sigma^2 T}\right) dx.
\end{aligned}
$$

Note that this calculation requires stock prices to be strictly positive.

The parameter μ is called the *expected return*, and σ the *volatility*. This model assumes μ and σ to be constant in time, which makes the analysis much simpler. Constant parameter models, however, are often poor reflections of reality.

Why should we believe the log-normal hypothesis about stock prices? Perhaps it would be more credible to suppose that the daily (or hourly or minute-by-minute) return is determined by a random event which we can model by flipping a coin. The log-normal model is the limit of such dynamics, as the time-frequency of the coin-flips tends to zero.

3.4.1 Log-normal dynamics as the limit of a binomial tree

Suppose then, we play a game in which at each stage is equivalent to flipping a coin. If the coin comes up heads our opponent pays us δx dollars and if it comes up tails then we pay δx dollars. Let X_i be the change in our fortune over the ith play. Each random variable in the sequence $\{X_i;\ i = 1, 2, \dots\}$ is independent and identically

distributed. The most general case is

$$X_i = \begin{cases} \delta x, & \text{w.p. } p; \\ -\delta x, & \text{w.p. } 1 - p. \end{cases}$$

Let $S_n = S_0 + \sum_{i=1}^n X_i$ be the random variable denoting our (possibly negative) fortune after the nth the coin toss. The sequence $\{S_n;\ n = 1, 2, \dots\}$ is a stochastic process known as a *random walk*. The path $S_1, S_2, \dots, S_n$ is chosen from a set of sample paths which form a multiperiod binary tree with

$$s_u = s_{\text{now}} + \delta x,$$
$$s_d = s_{\text{now}} - \delta x.$$

We would like to think of the stock price process as an infinitesimal random walk. In terms of our gambling game, the time interval between plays is δt and the stake is δx. To enable us to take the limit as these quantities tend to zero we recall a theorem of fundamental importance in probability.

Theorem 4 (Central Limit Theorem). *Let $X_1, X_2, \dots$ be a sequence of independent and identically distributed random variables with finite means μ and finite nonzero variances σ^2 and let $S_n = X_1 + X_2 + \dots + X_n$. Then*

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

converges in distribution to a Gaussian random variable with mean 0 and variance 1 as $n \to \infty$.

A Gaussian random variable with zero mean and unit variance is called *standard normal*.

The coin can be biased or fair, to keep things simple let's concentrate on the fair case; $p = 1 - p = \frac{1}{2}$. Then we have $\mu = \mathbb{E}(X_i) = 0$ and $\sigma^2 = var(X_i) = (\delta x)^2$. Thus

$$\Pr\left(\frac{S_n}{\sqrt{n}\delta x} \le x\right) \to \int_{-\infty}^x \frac{1}{\sqrt{2\pi}\delta x} e^{-\frac{y^2}{2(\delta x)^2}} dy, \quad \text{as } n \to \infty.$$

A simple random walk, however, won't adequately model a stock price process. Its fatal flaw is that it can wander onto the negative axis. Additionally, the rate of change in stock price with respect to time should be proportional to the price. For these reasons we need to re-think our approach.

Recall the way in which we modelled interest growth in discrete time. A bond worth $M(0)$ dollars now was worth $M(t) = B_t M(0)$ dollars a period t in the future, where $B_t = e^{rt}$. The quantity that's constant is not $\frac{dM}{dt}$ but rather the interest rate $r = \frac{1}{M}\frac{dM}{dt} = \frac{d}{dt}\log M$. Perhaps we should aim for a similar model for stock growth. Suppose, for instance, that

$$\frac{d}{dt}\log S_t = h_i,$$

where h_i is constant throughout the time interval $[(i-1)\delta t, i\delta t]$. Note that this certainly satisfies the criteria $\frac{dS_t}{dt} \propto S_t$. Solving this differential equation within each time tick leads us to

$$S_i = S_{i-1}e^{h_i\delta t},$$

or more generally

$$S_i = e^{h_i\delta t + h_{i-1}\delta t + \cdots + h_{j+1}\delta t}S_j, \quad \forall j < i, \qquad (3.8)$$

which is similar to the interest rate model. Unlike the interest rate model though, future stock prices are not known in advance, therefore each h_i should be a random variable. Suggested earlier was the idea of a random walk. Suppose then that each h_i is independently and identically distributed with

$$h_i = \begin{cases} \mu + \sigma/\sqrt{\delta t}, & \text{w.p. } 1/2, \\ \mu - \sigma/\sqrt{\delta t}, & \text{w.p. } 1/2. \end{cases}$$

The μ-part of the expression can be regarded as modelling long-term market trends; in compensation for our $p = 1 - p = \frac{1}{2}$ simplification. The σ-part is really where the randomness arises.

Note that this is the Cox-Ross-Rubinstein binomial tree model with

$$s_u = s_{\text{now}} e^{\mu\delta t + \sigma\sqrt{\delta t}},$$
$$s_d = s_{\text{now}} e^{\mu\delta t - \sigma\sqrt{\delta t}}.$$

Consider any time t. What is the probability distribution of stock prices at time t? Let's assume for simplicity that t is a multiple of δt, specifically $t = n\,\delta t$. From Equation 3.8 we have

$$S_n = S_0 e^{\sum_{i=1}^{n} h_i \delta t},$$

$$\log\left(\frac{S_n}{S_0}\right) = \sum_{i=1}^{n} h_i \delta t =: W_n, \quad \text{say.} \tag{3.9}$$

Now W_n is the sum of sequence of independently and identically distributed random variables with means $\mu\,\delta t$ and variances $\sigma^2 \delta t$ (check this). The Central Limit Theorem tells us that as $n \to \infty$, $\frac{W_n - n\mu\delta t}{\sqrt{n\,\delta t}\,\sigma} \to Z$ where Z is standard normal. Asymptotically then, as $\delta t \to 0$

$$\log\left(\frac{S_n}{S_0}\right) \to \mu\,n\,\delta t + \sigma\sqrt{n\,\delta t}\,Z,$$
$$\log\left(\frac{S_t}{S_0}\right) = \mu\,t + \sigma\sqrt{t}\,Z, \quad \text{since } t = n\,\delta t,$$
$$S_t = S_0\,e^{\mu t + \sigma\sqrt{t}\,Z}.$$

3.4.2 Implication for pricing options

We attached subjective transition probabilities (always equal to $1/2$) to our binomial tree because we wanted to recognise log-normal dynamics as the limit of a coin-flipping process. Now let's consider a single binomial branch corresponding to some specific time interval $[(i-1)\delta t, i\,\delta t]$ with δt near zero and use it to price options.

The structure of the tree is important. From Sections 3.1 and 3.2 we know the future values s_u and s_d are of particular relevance. Also in those sections we established that the transition probabilities p and $1 - p$ can be disregarded in favour of the risk-neutral probabilities $q = \frac{e^{r\delta t}S_{i-1} - s_d}{s_u - s_d}$ (returning momentarily to the case $r_i = r, \forall i$). Lastly, recall that we have the risk-neutral valuation formula for the price of an option with payoff V_T at maturity T:

$$V_{n-1}(\mathcal{F}_{n-1}) = B_{n-1}\mathbb{E}_{\mathbb{Q}}(B_T^{-1}V_T|\mathcal{F}_{n-1}).$$

In the current context

$$q = \frac{e^{r\delta t} - e^{\mu\delta t - \sigma\sqrt{\delta t}}}{e^{\mu\delta t + \sigma\sqrt{\delta t}} - e^{\mu\delta t - \sigma\sqrt{\delta t}}}.$$

Using the Taylor expansion of e^x in the neighbourhood of $x = 0$ we have (check this)

$$q = \frac{1}{2}\left(1 - \sqrt{\delta t}\frac{\mu - r + \sigma^2/2}{\sigma}\right) + o(\delta t).$$

where $o(h)$ represents a function of h such that $\frac{o(h)}{h} \to 0$ as $h \to 0$.

Our task now is clear. All we have to do is find the distribution of final values S_T when one uses a q-biased coin, then take the expected value of the payoff function V_T with respect to this distribution.

Proceeding in a similar fashion as before, let

$$\log\left(\frac{S_n}{S_0}\right) = \sum_{i=1}^{n} h_i \delta t =: W_n,$$

but in contrast to Equation (3.9), here the h_is are defined as

$$h_i = \begin{cases} \mu + \sigma/\sqrt{\delta t}, & \text{w.p. } q, \\ \mu - \sigma/\sqrt{\delta t}, & \text{w.p. } 1 - q. \end{cases}$$

It is straight forward, though a little messy, to verify that W_n is the sum of independent and identically distributed random variables with

$$\begin{aligned} \mathbb{E}_{\mathbb{Q}}(h_i\delta t) &= (r - \frac{1}{2}\sigma^2)\delta t + o(\delta t), \\ var_{\mathbb{Q}}(h_i\delta t) &= \sigma^2\delta t + o(\delta t). \end{aligned}$$

The Central Limit Theorem tells us the limiting distribution is Gaussian, and the preceding calculations tell us its mean and variance. Applying the Central Limit Theorem,

$$\frac{W_n - n(r - \frac{1}{2}\sigma^2)\delta t}{\sqrt{n\sigma^2\delta t}}$$

converges in distribution to a standard normal random variable Z. Therefore

$$\log\left(\frac{S_t}{S_0}\right) \rightarrow \left(r - \frac{1}{2}\sigma^2\right)t + \sigma\sqrt{t}Z$$

in distribution. To summarise, as $\delta t \rightarrow 0$, when using the biased coin associated with the risk-neutral probabilities,

$$S_t \rightarrow S_0 \exp\left[\left(r - \frac{1}{2}\sigma^2\right)t + \sigma\sqrt{t}Z\right],$$

where Z is normal with mean 0 and variance 1.

Remarkably, the statistical distribution of S_t under the measure $\mathbb{Q}$ depends on σ and r but not μ! It follows that for pricing derivative securities the value μ isn't really needed. So we may choose μ any way we please - there's no reason to require that it match the actual expected return of the stock under consideration. For instance, we may choose μ so that $\mu - r + \frac{1}{2}\sigma^2 = 0$ to simplify the calculations.

Now that we have the probability distribution of S_t in terms of Z we can apply the risk-neutral valuation formula to calculate the time-zero value of a claim. If C_T is the call option maturing at date T, struck at K, with $C_T = (S_T - K)^+$, then it's worth at time zero is

$$\begin{aligned}
\mathbb{E}_\mathbb{Q}(B_T^{-1}C_T) &= \mathbb{E}_\mathbb{Q}\left((S_0 e^{-\frac{1}{2}\sigma^2 T + \sigma\sqrt{T}Z} - Ke^{-rT})^+\right), \\
\\
&= \int_{\log(K/S_0)}^{\infty} (S_0 e^x - K)\frac{1}{\sigma\sqrt{2\pi T}}e^{\frac{-\left(x - \left(r - \frac{\sigma^2}{2}\right)T\right)^2}{2\sigma^2 T}}dx, \qquad (3.10) \\
\\
&= S_0\Phi(d_1) - Ke^{-rT}\Phi(d_2). \qquad (3.11)
\end{aligned}$$

where $\Phi(x) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{x} e^{-u^2/2}du$ is the cumulative distribution function of a standard normal random variable, and d_1 and d_2 are simplifying variables defined as

$$\begin{aligned}
d_1 &= \frac{1}{\sigma\sqrt{T}}\left[\log\left(\frac{S_0}{K}\right) + \left(r + \frac{1}{2}\sigma^2\right)T\right], \\
\\
d_2 &= \frac{1}{\sigma\sqrt{T}}\left[\log\left(\frac{S_0}{K}\right) + \left(r - \frac{1}{2}\sigma^2\right)T\right].
\end{aligned}$$

There is quite a bit of algebra involved in getting from Equation (3.10) to Equation (3.11), it's a good exercise. Typical manipulations to perform include "completing the square" and "changing of dummy variables".

A similar result may be obtained if P_T is the payoff of a European put option with strike K and expiry date T. In this case the analogue of Equation (3.11) is

$$\mathbb{E}_{\mathbb{Q}}\left(B_T^{-1}P_T\right) = Ke^{-rT}\Phi\left(-d_2\right) - S_0\Phi\left(-d_1\right). \tag{3.12}$$

Equation (3.12) ((3.11)) is Black-Scholes formula for the risk-neutral valuation of a European put (call) option.
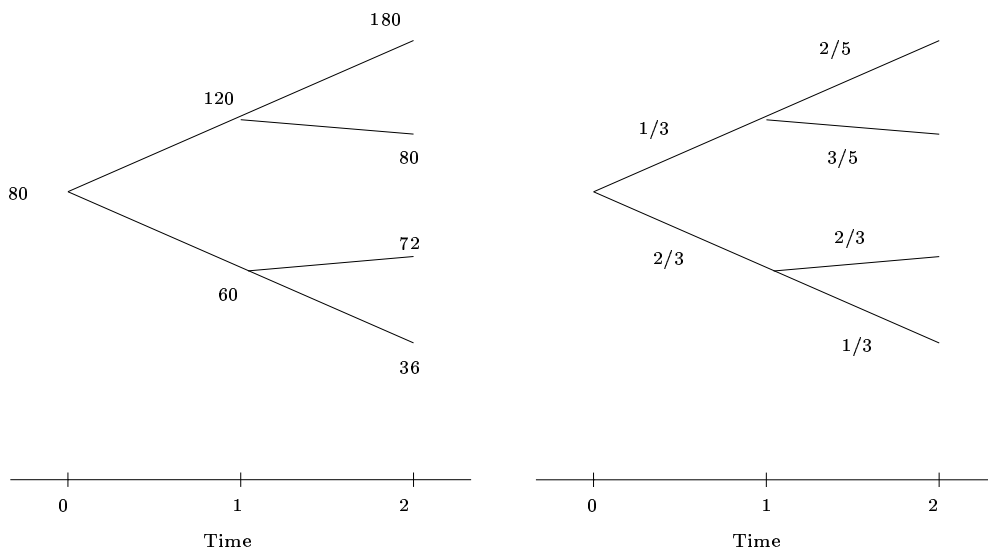
3.5   Appendix

This section contains a very brief review of the probabilistic ideas, such as measures, conditional expectation, filtrations, and martingales, which are relevant to financial calculus. At the level of this course, some of the best books on probability theory and stochastic processes include Grimmett & Stirzaker (1997), and Ross (1996). Dudley (1989) contains a useful collection of results, while any book with a title such as "Martingale methods in financial modelling", Musiela & Rutkowski (1997) is bound to be helpful.

In discrete time the stock price process is really just a sequence of random variables S_0, S_1, S_2, ..., S_T, or *stochastic process*. Given a measure $\mathbb{P}$ we have calculated the expected value of S_T, $\mathbb{E}_{\mathbb{P}}(S_T)$. The conditional expectation operator takes another parameter; the process history. For any stochastic process, the conditional expectation written

$$\mathbb{E}_{\mathbb{P}}(S_T|S_j, S_{j-1}, \ldots, S_0),$$

is the expected value of S_T given the process history up to some earlier time j. The history $\{S_j, S_{j-1}, \ldots, S_0\}$ is sometimes written $\mathcal{F}_j$. $\mathcal{F}_j$ is called a *filtration* and denotes the set (σ-algebra) of all decidable events given the history $\{S_j, S_{j-1}, \ldots, S_0\}$.

A stochastic process $\{S_n, n \geq 1\}$ with $\mathbb{E}_{\mathbb{P}}(|S_n|) < \infty$ for $n = 1, \ldots$ is said to be a

Figure 3.3: Stock process and martingale measure $\mathbb{P}$

$\mathbb{P}$-*martingale* process if

$$\mathbb{E}_{\mathbb{P}}(S_T | S_j, \dots, S_0) = S_j, \quad j = 0, \dots, N-1.$$

Grimmett & Stirzaker (1997) and also Ross (1996) both have excellent chapters on martingales.

The stock price process in Example 6 is a martingale with respect to the measure $\mathbb{Q}$ which assigns equal probabilities for steps up and down. Both measures operate on the same state space. In fact, every event which is possible under $\mathbb{Q}$ also has positive probability under $\mathbb{P}$, and vice versa. Any two measures which share this property are said to be *equivalent*.

Example 10. *The right pane of Figure 3.3 gives the martingale measure for the stock process in the left pane. It is easy to check that this is the case since*

$$\mathbb{E}_{\mathbb{P}}(S_1 | S_0 = 80) = 120\frac{1}{3} + 60\frac{2}{3} = S_0,$$

$$\mathbb{E}_{\mathbb{P}}(S_2 | S_1, S_0 = 80) = \left\{ \begin{array}{ll} 180\frac{2}{5} + 80\frac{3}{5}, & \textit{if } S_1 = 120 \text{ ;} \\ 72\frac{2}{3} + 36\frac{1}{3}, & \textit{if } S_1 = 60 \text{ .} \end{array} \right\} = S_1,$$

$$\mathbb{E}_{\mathbb{P}}(S_2|S_0 = 80) = \frac{1}{3}\mathbb{E}_{\mathbb{P}}(S_2|S_1 = 120, S_0 = 80) + \frac{2}{3}\mathbb{E}_{\mathbb{P}}(S_2|S_1 = 60, S_0 = 80) = S_0.$$

Theorem 3 refers to the discounted stock price process being a martingale under $\mathbb{Q}$. That is to say $\mathbb{Q}$ is such that

$$\mathbb{E}_{\mathbb{Q}}(B_j^{-1}S_j|S_0, S_1, \ldots, S_i) = B_i^{-1}S_i, \qquad \forall i \leq j. \tag{3.13}$$

It is sometimes convenient to use the change of variables $\tilde{S}_n = \frac{S_n}{B_n}$ this makes this premise (3.13)

$$\mathbb{E}_{\mathbb{Q}}(\tilde{S}_j|\tilde{S}_0, \tilde{S}_1, \ldots, \tilde{S}_i) = \tilde{S}_i, \qquad \forall i \leq j,$$

which appears simpler and more natural from the viewpoint of the martingale definition.

3.6 Exercises

1. How would the roulette examples change if the table was biased such that $\Pr(oddnumber) = \frac{3}{4} = 1 - \Pr(notanodd)$.

2. Derive the single step price for a European put option.

3. Show that the initial value of a contract with payoff function

$$V_T = \begin{cases} f_u, & \text{if the stock rises;} \\ f_d, & \text{otherwise,} \end{cases}$$

priced using the single step binomial model is

$$e^{rT}[qf_u + (1-q)f_d],$$

where $q = \frac{e^{rT}s_0 - s_d}{s_u - s_d}$.

4. Using the single step model show that the hedging strategy which will replicate a contract with payoff function

$$
V_T = \begin{cases} f_u, & \text{if the stock rises;} \\ f_d, & \text{otherwise,} \end{cases}
$$

   is given by contains ϕ units of stock, and ψ cash, where

$$
\begin{aligned}
\phi &= \frac{f_u - f_d}{s_u - s_d}, \\
\psi &= e^{-rT}(V_0 - \phi s_0).
\end{aligned}
$$

5. Evaluate $\mathbb{E}_\mathbb{Q}\left(B_T^{-1}(S_T - K)\right)$. Hence prove that the strike of a forward is fixed at $K = B_T S_0$ as claimed in Example 9.

6. Example 7 describes the replicating strategy to be followed should the stock price process follow a down, up, up path.

   (a) What would the strategy be if the stock went up, down, up?

   (b) What would the strategy be if the stock went up, down, down?

7. Using the binomial tree model of Figure 3.1, fill in a similar tree with the values of a European put option

   (a) with strike price $110.

   (b) with strike price $90.

8. Using the binomial tree model of Figure 3.1, fill in a similar tree with the values of a contract with payoff $|S_T - 100|$.

9. Using the log-normal model derive the value of a European put option T time units before expiry; Equation (3.12).

10. Show that Equations (3.11) and (3.12) satisfy the put-call parity.

11. Using the Cox-Ross-Rubinstein model for stock price movements, verify that the formula

$$C_{T-m} = e^{-rT} \sum_{i:u^i d^{m-i} s_{T-m} > K} \binom{m}{i} \left(\frac{e^{r\delta t} - d}{u - d} \right)^i \left(\frac{u - e^{r\delta t}}{u - d} \right)^{m-i} \left(u^i d^{m-i} s_{T-m} - K \right),$$

gives the value of a European call option struck at K with m time ticks until maturity.