4 CHAPTER

Asset Price Dynamics

4.0 INTRODUCTION

In Chapter 3 we derived upper and lower bounds for option prices using simple arbitrage arguments. Although these bounds limit the price of the option, they can be quite large. For example, consider a European call option with a strike price of 100, maturity date in six months, and an underlying asset price of 100. Let the interest rate be 6 percent. Using Results 3 and 4 in Chapter 3, we know that the value of the option must be less than 100 or greater than 2.96. This leaves us with lots of room in which to maneuver. To price options more precisely, we must make additional assumptions about the probability distribution describing the possible price changes in the underlying asset.

The purpose of this chapter is to study a model for the evolution of asset prices. The model needs to be simple enough to facilitate analysis, but complex enough to provide a reasonable approximation to the actual evolution of asset price movements. The model selected for presentation, with these characteristics in mind, is the lognormal distribution model, the "workhorse" for the subsequent options/futures pricing theory. It underlies the Black-Scholes model for pricing equity derivatives (Chapter 8) as well as the pricing of foreign currency derivatives (Chapter 11) and the special cases of the Heath-Jarrow-Morton model studied in Chapters 16 and 17. The following section provides a complete analysis of the lognormal distribution and justifies its selection as the basic model for asset price dynamics.

The lognormal distribution is well suited for continuous trading models and the use of calculus. Nonetheless, it is our experience that continuous trading models (and the use of calculus) is less intuitive than discrete trading models (and the use of algebra). For this reason we also introduce the binomial model.

The binomial model is cast in discrete time, and it is a very useful teaching tool for understanding the pricing and hedging of options/futures. Furthermore, if the binomial model is carefully constructed it can also serve as an approximation to the lognormal distribution and is useful in practice. In fact, in applications such as American option valuation, the binomial approximation to the lognormal distribution is the model of choice for many financial institutions.

Because of its simplicity, we will utilize the binomial model in this text to explain the arguments underlying the various options/futures pricing theories. Nonetheless, the lognormal distribution will always be lurking in the background, motivating and calibrating the models used in the various applications.

To fix discussion, the asset under consideration in this chapter will be called a stock. The analysis, however, applies equally well to most other assets and commodities addressed in this text.

4.1 THE LOGNORMAL DISTRIBUTION

A lognormal distribution for stock price returns is the standard model used in financial economics. Why? The answer is the topic of this section.

We show that given some reasonable assumptions about the random behavior of stock returns, a lognormal distribution is implied. These assumptions, in fact, characterize the lognormal distribution in a very intuitive manner. This intuition is important for our understanding because, to reiterate, the lognormal distribution is the "workhorse" for the subsequent derivative securities theory.

To motivate the analysis, consider a typical stock price chart as illustrated in Figure 4.1. A stock price evolution is usually very jagged, with peaks and valleys, sometimes separated by trend-like rises or declines. As all stock analysts and portfolio managers know, the future price of a stock is uncertain and very difficult to predict. For illustrative purposes, we have subdivided the time horizon $[0, T]$ into n equally spaced intervals of length Δ. By understanding the stock price process over each interval, we can understand it over the total horizon.

Figure 4.1 *A Typical Stock Price Chart*

We start by developing some notation. Let $S(t)$ be the stock's price at date t. Define z_t to be the continuously compounded return on the stock over the time interval $[t - \Delta, t]$, that is,

$$S(t) = S(t - \Delta)e^{z_t}. \tag{4.1}$$

To analyze the stock price process over the horizon $[0, T]$, we divide the period into n intervals of length Δ, when $T = n\Delta$. The stock price at the end of the first interval is denoted by $S(\Delta)$, at the end of the second interval by $S(2\Delta)$ and so forth, and at the end of the nth interval by $S(T)$. We can write the stock price $S(T)$ as the product of the ratios of the intervening stock prices:

$$S(T) = \left[\frac{S(T)}{S(T-\Delta)}\right]\left[\frac{S(T-\Delta}{S(T-2\Delta)}\right]\cdots\left[\frac{S(2)}{S(\Delta)}\right]\left[\frac{S(\Delta)}{S(0)}\right]S(0). \tag{4.2}$$

This follows because on the right side of this expression, we are always multiplying and dividing by the same stock prices. Next, substituting the definition of the continuously compounded return into Expression (4.2) gives the desired result:

$$S(T) = S(0)e^{z_\Delta + z_{2\Delta} + \dots + z_{T-\Delta} + z_T}. \tag{4.3}$$

To define:

$$Z(T) = z_\Delta + z_{2\Delta} + \dots + z_{T-\Delta} + z_T. \tag{4.4}$$

$Z(T)$ represents the continuously compounded return on the stock over the horizon $[0, T]$.[1]

$Z(T) = \ln[S(T)/S(0)]$ is seen to be the sum of the continuously compounded returns over the n intervals. This simple linear relationship is the reason for working with continuously compounded returns rather than with discretely compounded returns. For discretely compounded returns, a much more complex relationship applies.

To obtain a lognormal distribution for stock prices, we now impose some conditions on the probability distributions for the continuously compounded returns z_t.

These conditions are motivated by the empirical evidence. First, stock returns over successive intervals have been observed to be approximately statistically independent of each other. Second, over each interval, stock returns appear to be generated by the same distribution. Formally, we impose two assumptions:

Assumption A1. The returns $\{z_t\}$ are independently distributed.

Assumption A2. The returns $\{z_t\}$ are identically distributed.

The first assumption, A1, implies that the return over the interval $[t - \Delta, t]$, z_t, is of no use in predicting the return $z_{t+\Delta}$ over the next interval.

[1] This can be seen by substituting Expression (4.4) into (4.3) to obtain $S(T) = S(0)e^{Z(T)}$.

The second assumption, A2, implies that the return z_t does not depend upon the previous stock price $S(t - \Delta)$.

These two assumptions together imply that stock prices follow a **random walk**.[2] This characteristic of stock prices has often been associated with "efficient market" theory.[3]

Given these two assumptions, we now describe how the return changes as the size of the time interval Δ declines. We want to ensure that the characteristics of the return over each interval, as described in Assumptions A1 and A2, remain intact as the size of the time interval becomes smaller. To achieve this end, we add two more assumptions:

Assumption A3. The expected continuously compounded return can be written in the form

$$E[z_t] = \mu\Delta,$$

where μ is the expected continuously compounded return per unit time.

Assumption A4. The variance of the continuously compounded return can be written in the form

$$\text{var}[z_t] = \sigma^2\Delta,$$

where σ^2 is the variance of the continuously compounded return per unit time.

Assumption A3 states that the expected value of the continuously compounded return equals a constant μ times the length of the interval Δ.

Assumption A4 states that the variance of the continuously compounded return equals a constant σ^2 times the length of the interval Δ.

Both the expected return and the variance of the return are seen to be proportional to the length of the time interval. Thus, as the length of the time interval decreases, these two moments of the stock return's distribution decrease proportionately.

Technically, these assumptions ensure that as the time interval decreases, the behavior of the distribution for $Z(T)$ does not explode nor degenerate to a fixed point. It remains random and similar in appearance to any other size interval, appropriately magnified.

Given these four assumptions, the expected continuously compounded return over the horizon $[0, T]$ is

$$\begin{aligned} E[Z(T)] &= E(z_\Delta) + E(z_{2\Delta}) + \dots + E(z_t) \\ &= \sum_{j=1}^{n} \mu\Delta \qquad \text{using Assumptions A2 and A3} \\ &= \mu T. \end{aligned} \tag{4.5}$$

[2] A random walk does not imply any particular probability distribution for stock price changes.
[3] The efficient market theory is described in Fama (1970, 1991).

The variance of the continuously compounded return over the horizon $[0, T]$ is

$$\begin{aligned} \text{var}[Z(T)] &= \text{var}(z_\Delta) + \text{var}(z_{2\Delta}) + \dots + \text{var}(z_T) \\ &= \sum_{j=1}^{n} \sigma^2\Delta \qquad \text{using Assumptions A2 and A4} \\ &= \sigma^2 T. \end{aligned} \tag{4.6}$$

At this point it may appear as if we have not made any restrictive assumptions about the probability distribution for each continuously compounded return z_t and thus $Z(T)$. But, Assumptions A1 to A4 are quite powerful and imply that for infinitesimal time intervals, the distribution for the continuously compounded return z_t has a normal distribution with mean $\mu\Delta$ and variance $\sigma^2\Delta$. The proof of this result relies on the Central Limit Theorem from probability theory. Because the proof provides no additional insight, we leave it to interested readers to pursue the proof in the references (see Cox and Miller, 1990). This result, in turn, can be shown to imply that stock prices are lognormally distributed.

Let us summarize what we have achieved. Given a horizon $[0, T]$, we divided it into n intervals of length Δ and examined the distribution of the continuously compounded returns over each interval. We imposed Assumptions A1 through A4 on the nature of these returns based on empirical considerations. The assumptions implied that for infinitesimal intervals, returns are normally distributed. Since the sum of n independent normally distributed random variables is itself normally distributed, using Expression (4.4), we see that $Z(T) = \ln[S(T)/S(0)]$ is normally distributed with mean μT-expression (4.5) and variance $\sigma^2 T$-expression (4.6). But, this is equivalent to stating that the stock price $S(T)$ is *lognormally distributed.*

In fact, Assumptions A1 through A4 characterize the lognormal distribution for stock returns. Why? If $Z(T)$ has a normal distribution, z_t will satisfy Assumptions A1 through A4 as well. Thus we have obtained our "workhorse" model for the evolution of stock prices, as exhibited in Figure 4.1.

EXAMPLE **Lognormally Distributed Stock Prices**

Suppose that the expected return, expressed on a continuously compounded basis, is 15 percent per year and the volatility of the return is 25 percent per year. The distribution for the continuously compounded return over a two-year period is normally distributed with a mean of $15 \times 2 = 30$ percent and volatility of $25 \times \sqrt{2} = 35.36$ percent. ■

Figure 4.2 displays typical shapes for a normal and a lognormal distribution. Part A shows the distribution for the continuously compounded returns. Note that these returns can be negative, given that the normal distribution is defined for both positive

FIGURE 4.2 *Part A: Normal Distribution for Logarithm of the Price Relative, Z(T)*
Part B: Lognormal Distribution for S(T)

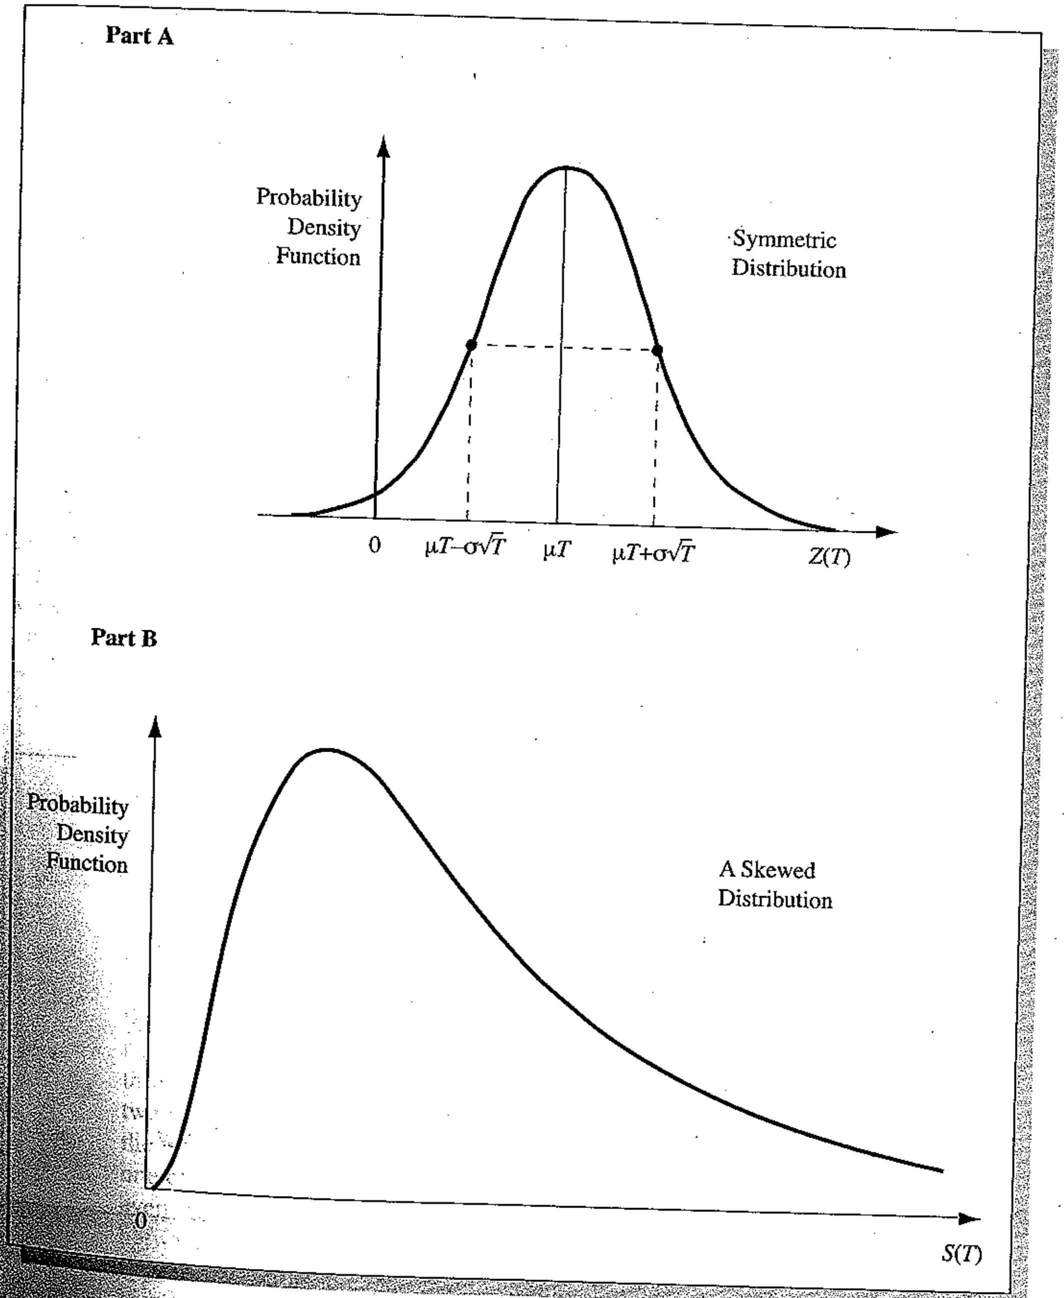

and negative values.[4] Part B shows the distribution for the stock price at date T. The lognormal distribution is only defined for positive values.[5]

If the stock price at date T is described by a lognormal distribution, the expected stock price at date T given today's price can be shown to be

$$E[S(T) \mid S(0)] = S(0)\exp(\mu T + \sigma^2 T/2). \tag{4.7}$$

A proof of this result is given in the chapter Appendix; we will use it in Chapter 5.

The assumption that stock prices are lognormally distributed is a convenient assumption and one that we will use extensively. It allows us to derive relatively simple expressions for different types of derivative securities. For example, this assumption is used in the Black-Scholes option model, which is the standard basic model for pricing equity options. However, simple convenience is not necessarily a sufficient justification for employing a particular assumption, and this is why we motivated this assumption based on four intuitive and empirically verifiable assumptions. This issue is briefly discussed again in Section 4.5.

4.2 THE BASIC IDEA (BINOMIAL PRICING)

Given that the lognormal distribution for stock price movements has been discussed, we now consider another related model, the binomial model. The binomial model is useful as a teaching tool in understanding options/futures pricing and hedging theory. We will subsequently relate this binomial model back to the lognormal distribution.

FIGURE 4.3 *Binomial Pricing*

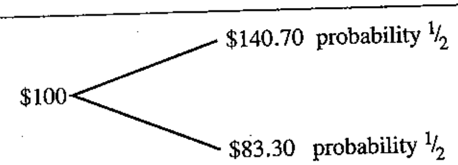

The return is defined as the price relative.
If the stock is \$140.70, the return is 1.407.
If the stock is \$83.30, the return is 0.833.

Expected return $= (1.407 + 0.833)(\tfrac{1}{2}) = 1.12.$

Variance of return $= [(1.407 - 1.12)^2 + (0.833 - 1.12)^2](\tfrac{1}{2})$
$= (0.0824 + 0.0824)(\tfrac{1}{2})$
$= 0.0824$

[4]The normal distribution is defined from minus infinity to plus infinity.
[5]The lognormal distribution is defined from zero to plus infinity.

FIGURE 4.4 *Multiperiod Extension*

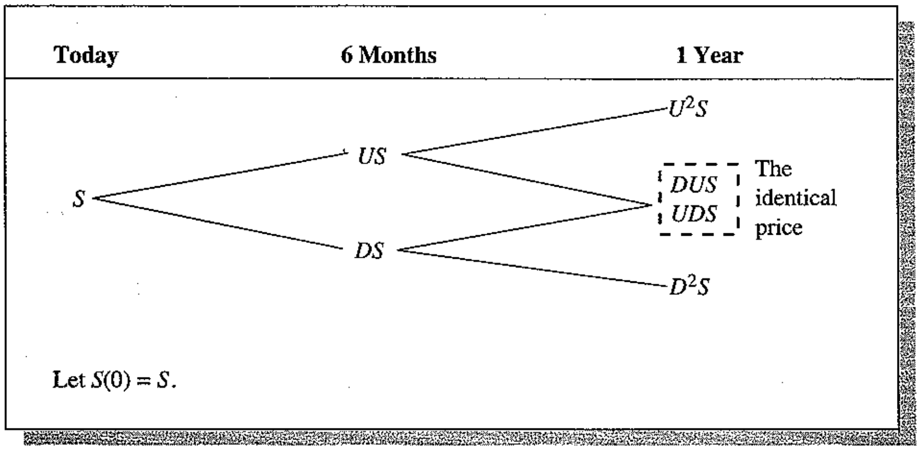

To be concrete, we must first consider a numerical example. Let the price of a stock today be \$100. We are interested in the stock's price in a year. For simplicity, assume that the stock does not pay any dividends over this period and that at the end of the year, the stock's price can take on only one of two possible values: either \$140.70 with probability $\frac{1}{2}$ or \$83.30 with probability $\frac{1}{2}$ (see Figure 4.3).

The expected dollar return on the stock is 1.12, where the **dollar return** is defined to be the price relative (one plus the percent return).

We can represent the stock price at the end of one year, $S(1)$, in the following manner:

$$S(1) = \begin{cases} U_0 S(0) & \text{if the stock price moves "up"} \\ D_0 S(0) & \text{if the stock price moves "down,"} \end{cases}$$

where $S(0)$ is the initial stock price, \$100. U_0 is called the **up-factor** with $U_0 = 1.407$, and D_0 is called the **down-factor** with $D_0 = 0.833$.

The assumption that the stock price can take only one of two possible values at the end of each interval is referred to as the **binomial model**. We could have alternatively assumed that at the end of each interval the stock price could have one of three (or more) possible values. The restriction to the binomial model is made for two reasons. The first is for simplicity. Jumping ahead, Figure 4.4 shows that even the binomial model gets complicated. Had we assumed that the stock could take one of three possible values at the end of each interval, then Figure 4.4 would be even more cluttered and complex. Second, for most purposes, as the time interval between price movements declines, this assumption is not as restrictive as it appears. We will show later that the binomial model can be used to approximate a lognormal distribution.

Regardless, you might argue that it is unrealistic to assume that the stock price can have only one of two possible values in a year. Given the multitude of events that may happen, one might expect a large number of possible values. We agree. As a first step toward accommodating this possibility, let us divide the one-year period into two subintervals of length six months. At the end of each of these six-month periods, it is now assumed again that the stock price can take on one of two possible values:

$$S(t+1) = \begin{cases} US(t) & \text{if the stock price moves "up"} \\ DS(t) & \text{if the stock price moves "down,"} \end{cases}$$

where $S(t)$ is the stock price at date t, and U and D are constants, U being greater than D.

The range of possible outcomes is shown in Figure 4.4. There are now three possible prices at the end of the year (U^2S, UDS, D^2S). Such a figure is referred to as a **lattice.**

We wish to make a number of comments about Figure 4.4. First, what are the values of U and D? A complete answer will be given in the next section. For the moment, notice that it is unreasonable to expect the initial values for the up and down factors, $U = 1.407$ and $D = 0.833$, to remain unchanged as the number of intervals increases. Because, if they remained unchanged, then as the number of periods increased the magnitude of the largest stock price would explode (approach infinity). To avoid this result, the size of the up and down factors must depend on the number of intervals.

Second, we assumed that the up and down factors do not depend on time and are state independent.[6] These assumptions imply that the lattice recombines at the end of each interval, giving a total of three distinct possible prices for the stock at year end. If we had n intervals per year, n being a positive integer, there would be $(n + 1)$ possible stock prices at year end. We could relax the assumptions of time and state independence, but it would result in an increase in the complexity of the lattice. An illustration of this complexity occurs when we talk about the pricing of interest-rate derivative securities in Chapter 15.

4.3 FORMAL DESCRIPTION (BINOMIAL PRICING)

In order to formalize[7] the description of the lattice, we divide the horizon $[0, T]$ into n periods of equal length Δ, where $T \equiv n\Delta$. Let $S(t)$ denote the stock price at date t, where $t = 0, \Delta, 2\Delta, \dots, n\Delta$. Remember that $S(0)$ denotes today's stock price. All future stock prices are uncertain.

At some intermediate date t, the stock price next period, date $t + \Delta$, would take the values

[6]By state independent, we mean that the up and down factors do not depend on the level of the stock price.

[7]This is not intended to be a rigorous derivation. We want to concentrate on the underlying economics without getting too involved in the mathematics. For more rigor, see the references.

FIGURE 4.5 *Multiperiod Bionomial Pricing*

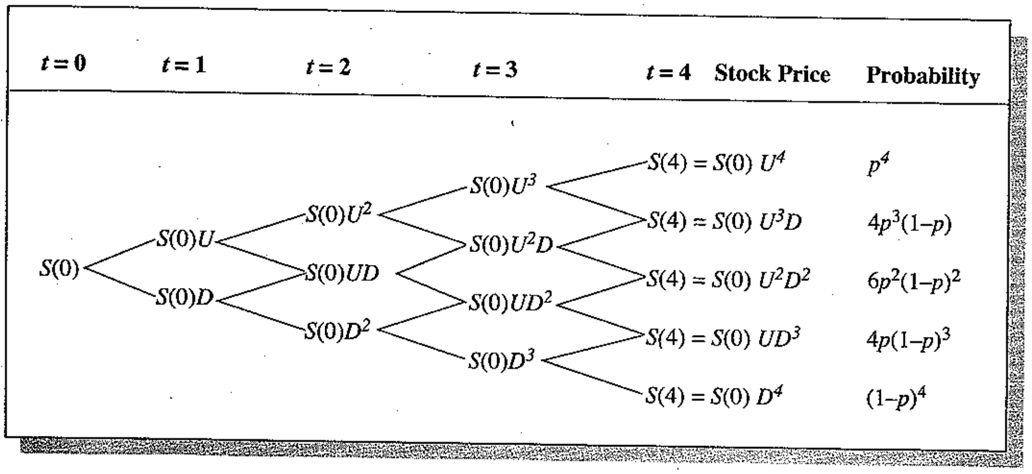

$$S(t + \Delta) = \begin{cases} S(t)U & \text{with probability } p \\ S(t)D & \text{with probability } 1 - p. \end{cases} \tag{4.8}$$

Taking $\Delta = 1$, a lattice of four intervals is shown in Figure 4.5. Observe that at the end of four intervals, the stock price $S(4)$ can have one of five possible values.

The probabilities are determined by considering the number of up and down transitions along all the feasible paths through the lattice. At date T, after n intervals, there are $(n + 1)$ possible values for the stock price $S(T)$. These values and their associated probabilities are listed in Table 4.1 on the next page. The probabilities are easily computed, and the resulting distribution for the time T stock price is known as a **multinomial distribution**. Tables for its values can be readily found in standard statistical software.

The formal description of the binomial model is now complete. To use this model in applications, the specification of the factors U and D is crucial. Different choices of U and D will generate different models for the stock price. Next we will show how to specify U and D so that the lattice of stock prices will approximate the lognormal distribution.

4.4 THE BINOMIAL APPROXIMATION TO THE LOGNORMAL DISTRIBUTION

We now show how to use the binomial model of the previous section to approximate the lognormal distribution of Section 4.1. It is done by choosing the up (U) and down (D) magnitudes in a clever fashion.

TABLE 4.1 *Stock Prices at Date T for a Lattice with n Intervals*

$S(T)$	PROBABILITY
$S(0)U^n$	p^n
$S(0)U^{n-1}D$	$\binom{n}{1}p^{n-1}(1-p)$
$S(0)U^{n-2}D^2$	$\binom{n}{2}p^{n-2}(1-p)^2$
	$\vdots$
$S(0)UD^{n-1}$	$\binom{n}{n-1}p(1-p)^{n-1}$
$S(0)D^n$	$(1-p)^n$

where the binomial coefficient $\binom{n}{k}$ is defined by

$$\binom{n}{k} = \frac{n \times (n-1) \times (n-2) \ldots 2 \times 1}{[k \times (k-1) \times \ldots \times 2 \times 1][(n-k) \times (n-k-1) \times \ldots \times 2 \times 1]}.$$

For example, when $n = 4$, the values for $k = 3$ and 2 are:

$$\binom{4}{3} = \frac{4 \times 3 \times 2 \times 1}{[3 \times 2 \times 1] \times 1} = 4 \quad \text{and} \quad \binom{4}{2} = \frac{4 \times 3 \times 2 \times 1}{[2 \times 1] \times [2 \times 1]} = 6.$$

Recall that the binomial representation assumes that, at the end of each interval, the stock's return can take only one of two possible values. Let us rewrite the binomial representation as

$$\ln[S(t)/S(t-\Delta)] \equiv z_t = \begin{cases} \mu\Delta + \sigma\sqrt{\Delta} & \text{with probability } 1/2 \\ \mu\Delta - \sigma\sqrt{\Delta} & \text{with probability } 1/2. \end{cases} \tag{4.9}$$

With probability $1/2$ the stock's (continuously compounded) return goes "up" to $\mu\Delta + \sigma\sqrt{\Delta}$, and with probability $1/2$ the stock's return goes "down" to $\mu\Delta - \sigma\sqrt{\Delta}$. The choice of the probability of an upward movement to be $1/2$ is justified subsequently.

The expected return over $[t-\Delta, t]$ is

$$\begin{aligned} E[z_t] &= (\mu\Delta + \sigma\sqrt{\Delta})(1/2) + (\mu\Delta - \sigma\sqrt{\Delta})(1/2) \\ &= \mu\Delta \end{aligned}$$

and the variance is

$$\begin{aligned}\text{var}[z_t] &= (\sigma\sqrt{\Delta})^2(1/2) + (-\sigma\sqrt{\Delta})^2(1/2)\\ &= \sigma^2\Delta.\end{aligned}$$

The expected return, $\mu\Delta$, is often called the **drift** because it is the value to which the stock return drifts before it is shocked by $+\sigma\sqrt{\Delta}$ or $-\sigma\sqrt{\Delta}$ (see Expression (4.9)). The square root of the term σ^2 is often called the stock's **volatility** (σ) because it reflects the size of the random shocks in the stock's return as it moves through time.

We now argue that Expression (4.9) approximates a lognormal distribution. Note that, by construction, Expression (4.9) satisfies Assumptions A1 through A4 in Section 4.1. First, z_t is independently and identically distributed since the probabilities (1/2), the drift μ, and the volatility σ do not change with t (Assumptions A1 and A2). Second, Assumptions A3 and A4 are seen to be satisfied by the expected return and variance of Expression (4.9). They are both proportional to the length of the time period Δ. Thus, by the argument used in Section 4.1, for infinitesimal intervals (as Δ tends to zero), z_t is approximately normally distributed.[8]

This is an important observation because it implies (as argued earlier) that the binomial representation in Expression (4.9) *approximates a lognormal distribution.*

Using Expression (4.9), the stock price at date t can be written in the form

$$S(t) = S(t-\Delta)\begin{cases}\exp(\mu\Delta + \sigma\sqrt{\Delta}) & \text{with probability } 1/2\\ \exp(\mu\Delta - \sigma\sqrt{\Delta}) & \text{with probability } 1/2.\end{cases} \tag{4.10}$$

Given this expression, we can easily identify the up and down factors in Figure 4.4 in terms of the instantaneous expected return per unit time, μ, the instantaneous volatility per unit date, σ, and the length of the interval, Δ. They are

$$U = \exp(\mu\Delta + \sigma\sqrt{\Delta})$$

and

$$D = \exp(\mu\Delta - \sigma\sqrt{\Delta}).$$

A final comment is in order. Because we are only interested in approximating a lognormal distribution as Δ gets small via a binomial representation, the representation of stock price movements in Expression (4.10) is not uniquely determined. There are other ways of representing stock price movements that satisfy Assumptions A1 through A4.[9]

---

[8] A formal proof is given in Cox and Miller (1990, Chapter 5).

[9] The representation used in Cox, Ross, and Rubinstein (1979) is

$$S(t) = S(t-\Delta)\begin{cases}\exp(\sigma\sqrt{\Delta}) & \text{with probability } [1 + (\mu/\sigma)\sqrt{\Delta}]/2\\ \exp(-\sigma\sqrt{\Delta}) & \text{with probability } [1 - (\mu/\sigma)\sqrt{\Delta}]/2.\end{cases}$$

EXAMPLE

Binomial Lattice Approximation to the Lognormal Distribution

Suppose that the expected return μ is 11 percent per year and the volatility σ is 25 percent per year. These numbers can be calibrated to market data.

In Figure 4.6 the horizon is one year. In Part A we have split this into two six-month intervals, implying the number of periods, $n = 2$, and the length of each interval, $\Delta = \frac{1}{2} = 0.5$. Therefore, the expected drift is

$$\mu\Delta = 0.11 \times .05 = 0.055$$

and the volatility over the interval is

$$\sigma\sqrt{\Delta} = 0.25 \times \sqrt{0.5} = 0.1768.$$

Hence, from Expression (4.9) the one-period continuously compounded returns can be written as

$$z_t = \begin{cases} 0.2318 & \text{with probability } \frac{1}{2} \\ -0.1218 & \text{with probability } \frac{1}{2}, \end{cases}$$

FIGURE 4.6 *Binomial Stock Price Lattices*

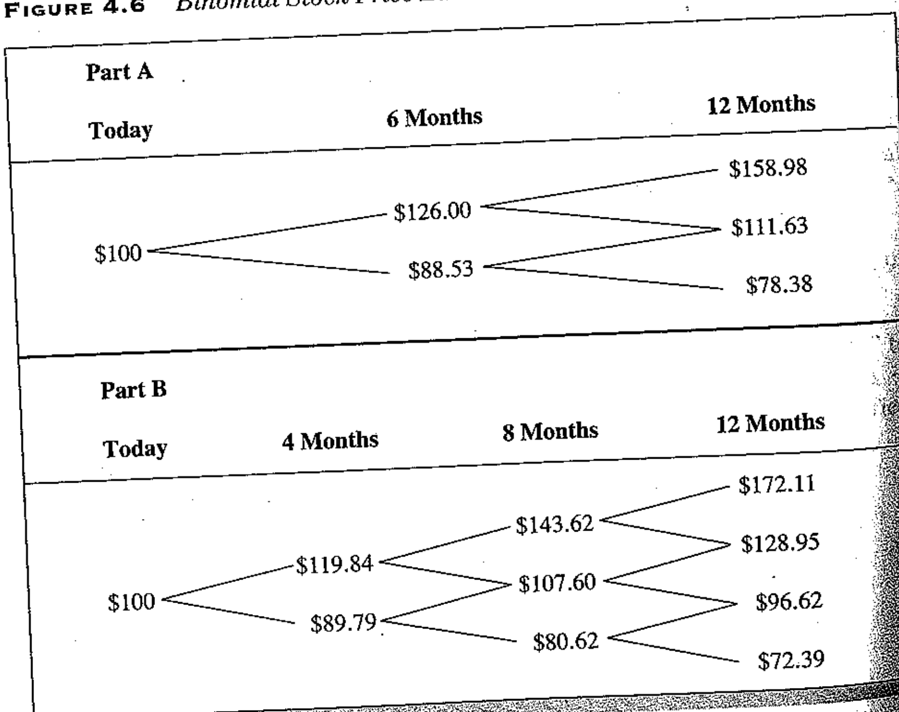

Table 4.2 *Size of the Up and Down Factors*

Number of Intervals	Length of Intervals	Up Factor	Down Factor
n	Δ	U	D
1	1	1.4333	0.8694
2	1/2	1.2609	0.8853
3	1/3	1.1984	0.8979

and from Expression (4.10), the binomial model is

$$S(t+\Delta) = S(t)\begin{cases} 1.2609 & \text{with probability } \tfrac{1}{2} \\ 0.8853 & \text{with probability } \tfrac{1}{2}. \end{cases}$$

In Figure 4.6, it is assumed that the initial stock price is \$100.

In Part B we have divided the twelve-month interval into three subperiods. As an exercise, check that you can reproduce the numbers in Part B. In this case, the binomial model is

$$S(t+\Delta) = S(t)\begin{cases} 1.1984 & \text{with probability } \tfrac{1}{2} \\ 0.8979 & \text{with probability } \tfrac{1}{2}. \end{cases}$$

Both these binomial models approximate the lognormal distribution, but the model with $n = 3$ does a better job than does the model with $n = 2$. The approximation improves as n increases. ■

Note two points from this example. The size of the up and down factors, U and D, change as the length of the interval changes. This relationship of U and D to the time interval Δ is summarized in Table 4.2. As the length of the interval Δ decreases, the size of the up factor U decreases toward 1 and the size of the down factor D increases toward 1. It is always the case that the size of the up factor U is greater than the size of the down factor D. This model will be utilized repeatedly later on in the text to understand the pricing and hedging of options and futures.

Example

Lognormal Approximation

We use Expression (4.10) to generate binomial stock price movements, which approximate a lognormal distribution. A random sequence of coin tosses is used to determine whether the stock price goes up or down each interval. If the current stock price is $S(t)$ and the coin toss is "heads," the stock price next interval is

$$S(t+\Delta) = S(t)\exp(\mu\Delta + \sigma\sqrt{\Delta}).$$

FIGURE 4.7 *Pattern of Price Movements*

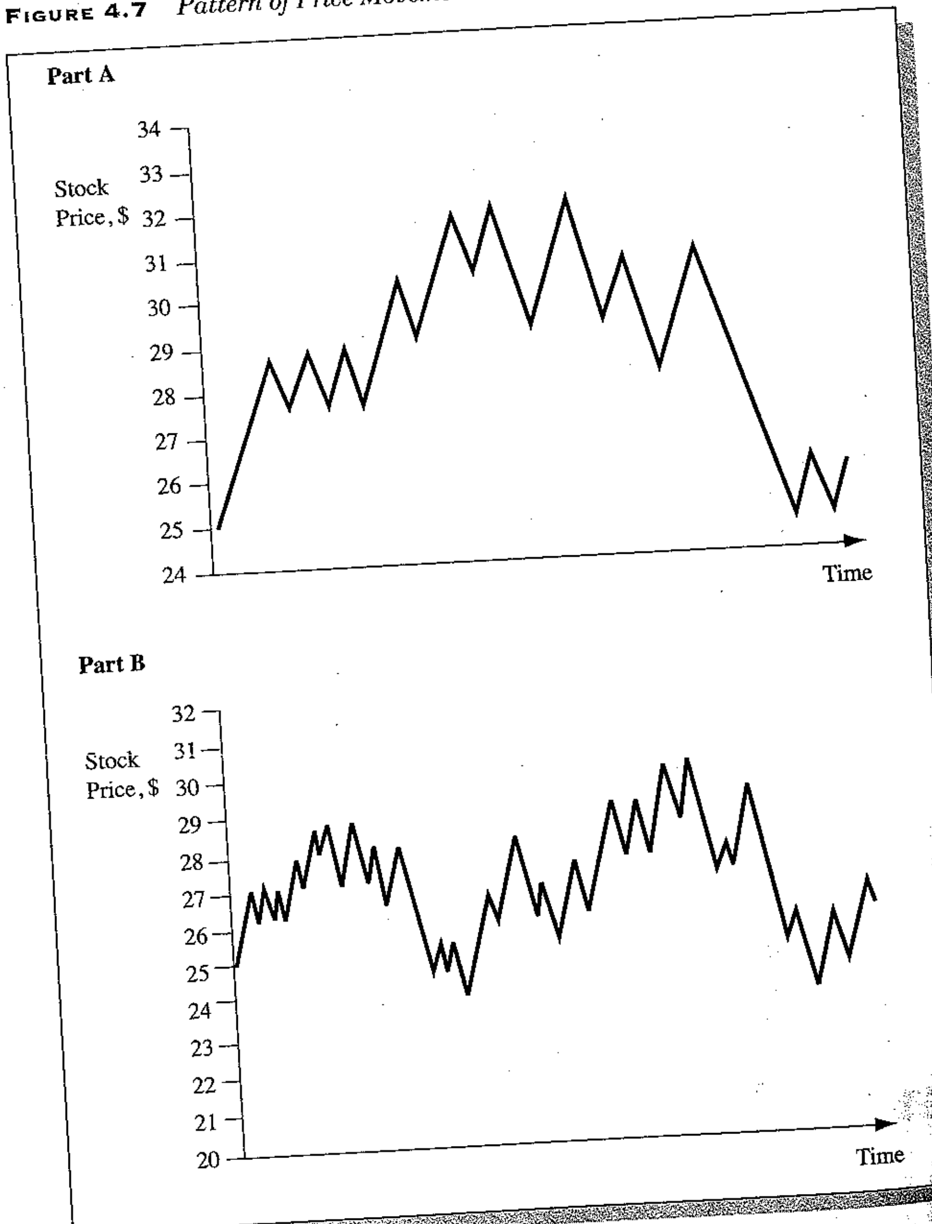

If the coin toss is "tails," the stock price next interval is

$$S(t + \Delta) = S(t)\exp(\mu\Delta - \sigma\sqrt{\Delta}).$$

In Figure 4.7, Part A, a three-day interval is used, so $\Delta = 3/365$. In Part B, a one-day interval is used, so $\Delta = 1/365$. You should compare these figures to Figure 4.1. ■

4.5 EXTENSIONS

We now discuss generalizations of the lognormally distributed stock price process. Recall that the lognormal distribution is characterized by Assumptions A1 through A4. Changing any of these assumptions will imply a different stock price distribution; the two assumptions most often modified are Assumptions A3 and A4. For example, the mean return μ and the variance of the return σ^2 can both be made functions of the stock price. This modification changes the stock price process to a non-lognormal distribution.

There are many reasonable distributions for which the return's variance depends upon the stock price level. For example, in some markets it is observed that the variance of price changes increases as the stock price increases. In terms of returns, it implies that the variance of the stock's return decreases as the stock price increases. We can incorporate this idea by assuming that

$$\text{var}[z_t] = \eta^2\Delta/S(t),$$

where η is the "new" volatility term.

This modification of Assumption A4 yields a **stochastic volatility model** for the stock price, which causes a number of complications. First, the lattice may not recombine, implying that in Figure 4.4 we would have four different price levels at the terminal time.[10] This condition can cause computing problems when the number of intervals is large. Second, successive price changes will no longer be independently distributed, which complicates the statistical procedures employed when estimating the parameters μ and η.

4.6 STOCHASTIC DIFFERENTIAL EQUATION REPRESENTATION

To read the academic literature on option pricing, one must be acquainted with the stochastic differential equation representation of lognormally distributed stock prices. We hereby provide a simple introduction.

[10] See Nelson and Ramaswamy (1990) for a description of how to use binomial processes to approximate different types of processes.

An elegant way to represent Assumptions A1 through A4 for continuous compounded returns is[11]

$$z_{t+\Delta} = \mu\Delta + \sigma[W(t + \Delta) - W(t)],$$

where $[W(t + \Delta) - W(t)]$ is a normally distributed random variable with zero mean and variance Δ.

The above equation is usually expressed in terms of stock price changes. Recall that, by definition,

$$z_{t+\Delta} = \ln[S(t + \Delta)/S(t)] = \ln S(t + \Delta) - \ln S(t).$$

Hence we can write Assumptions A1 through A4 alternatively as

$$\ln S(t + \Delta) - \ln S(t) = \mu\Delta + \sigma[W(t + \Delta) - W(t)]. \tag{4.11}$$

Replacing discrete changes with infinitesimal changes, that is, $dt \cong \Delta$, $d\ln S(t) \cong \ln S(t + \Delta) - \ln S(t)$, and $dW(t) \cong W(t + \Delta) - W(t)$, we obtain

$$d\ln S(t) = \mu dt + \sigma dW(t). \tag{4.12}$$

Expression (4.12) is the form by which Assumptions A1 through A4 most often appear in the literature.

For distributions more complex than the lognormal, Expression (4.12) can be generalized to

$$d\ln S(t) = \mu[t, S(t)]dt + \sigma[t, S(t)]dW(t), \tag{4.13}$$

where $d\ln S(t)$ represents the change in the natural logarithm of the stock price from date t to $t + dt$, with dt being an infinitesimal change in time. $\mu[t, S(t)]$ is the instantaneous expected return per unit time, and $dW(t)$ is a Brownian motion.

A **Brownian motion**, by definition, is a random variable that is normally distributed with zero mean, variance dt, and has independent increments (that is, $dW(t)$ and $dW(t + dt)$ are independently distributed).

Note that, in this general form, both the mean and volatility are functions of date t and the current stock price, $S(t)$. Expression (4.13) is called a differential equation

---

[11]The expected value of z_t is

$$E[z_t] = \mu\Delta$$

and the variance is

$$\begin{aligned} \text{var}[z_t] &= \sigma^2\text{var}[\Delta W(t)] \\ &= \sigma^2\Delta. \end{aligned}$$

because the stock price $S(t)$ is only defined implicitly by describing its changes through time.

Different assumptions about the form of the volatility give rise to different solutions $S(t)$ to this stochastic differential equation. The standard assumption is to assume that μ and σ are constants, which is the form of the equation implied by Expression (4.12). The solution for $S(t)$ in this case is a lognormal distribution. It is the stock price distribution underlying the Black-Scholes option pricing model and the special cases of the Heath-Jarrow-Morton model studied later.

4.7 COMPLICATIONS

We now consider various complications to the preceding theory.

Lognormal Distribution

If we look at the empirical distribution of continuously compounded returns, we find that the tails of the distribution are fatter than those expected by a normal distribution. This condition is inconsistent with a lognormal distribution for stock prices. By examining the dynamic properties of stock price changes, there is also some evidence that the volatility of the distribution changes,[12] which is inconsistent with a lognormal distribution for stock prices because stochastic volatility causes the distribution of returns to have fat tails. There is a growing literature using stochastic volatility stock price models to price derivative securities. We will return to this point shortly.

Continuous Trading

In the binomial model approximation to the lognormal, the length of the trading period Δ decreases in size as we increase the number of intervals. In the limit, the length of the trading period becomes infinitesimal and implies that trading is approximated as being continuous. This, of course, is not true because there are holidays and weekends on which markets close.

The closing of markets on weekends and holidays can cause Assumptions A1 through A4 to be violated. In particular, French (1980) has documented that over weekends the volatility of returns differs from that during the week, violating Assumption A4, which assumes a constant volatility. If we know the different volatilities during the different periods, we can introduce these complications into our description of stock price movements via the generalizations discussed in Section 4.5.

The tradeoff for increased realism will be increased complexity. The choice between the two (realism versus simplicity) is made with the use of the models in mind.

[12]See Schwert (1989) and Haugen, Talmor, and Torous (1991).

Trading rooms often prefer realism, while corporate treasury departments often prefer simplicity.

Continuously Changing Prices

Assumptions A3 and A4 imply that for infinitesimal time intervals, returns are normally distributed with a variance approaching zero. It implies that stock price changes will be quite small, and in the limit, continuously changing. Unfortunately, an institutional feature of most markets is the existence of minimum allowed price changes. For example, in equity markets the minimum price change is usually $^1/_{16}$ per share. We have ignored this institutional feature, and for most of the book we will continue to ignore it because, for large dollar positions, continuously changing prices is a reasonable approximation.

However, this institutional feature may affect how we estimate different parameters and test the models. For example, academic studies often ignore deep-out-of-the-money options near expiration because option prices are near zero, and $^1/_{16}$ can be a large percent of the option's value. Traders also recognize these difficulties and will often refrain from trading deep-out-of-the-money options near expiration.

4.8 SUMMARY

To price derivative securities, we need a way of representing the evolution of the future prices of an asset. We need a model that is simple enough to perform analysis but complex enough to provide a realistic approximation. The lognormal distribution is our selection.

However, to facilitate understanding we study the binomial model. We show how to specify the up and down factors such that as the number of intervals increases—or equivalently, the length of each interval decreases—the binomial model approximates a lognormal distribution. The binomial form of this representation will also be the model used to explain the pricing and hedging of equity, stock index, foreign currency, commodity, and interest rate derivatives in subsequent chapters.

REFERENCES

References for the material on convergence:

Cox, D. R., and H. D. Miller, 1990. *The Theory of Stochastic Processes*. London: Chapman Hall.

Cox, J., S. Ross, and M. Rubinstein, 1979. "Option Pricing: A Simplified Approach." *Journal of Financial Economics* 7, 229–264.

Hoel, G. H., S. C. Port, and C. H. Stone, 1971. *Introduction to Probability Theory*. Boston: Houghton Mifflin Company.

Nelson, D. B., and K. Ramaswamy, 1990. "Simple Binomial Approximations in Financial Models." *Review of Financial Studies* 3, 393–430.

References for properties of asset price distributions:

French, K. R., 1980. "Stock Returns and the Weekend Effect." *Journal of Financial Economics* 9, 55–69.

Haugen, R. A., E. Talmor, and W. N. Torous, 1991. "The Effect of Volatility Changes on the Level of Stock Prices and Subsequent Expected Returns." *Journal of Finance* 46, 985–1007.

Richardson, M., and T. Smith, 1993. "A Test for Multivariate Normality in Stock Returns." *Journal of Business* 66, 295–321.

Schwert, G. W., 1989. "Why Does Stock Market Volatility Change Over Time?" *Journal of Finance* 44, 1115–1153.

For a relatively simple explanation of the properties of the lognormal distribution, see:

Ingersoll, J. E., 1987. *Theory of Financial Decision Making*. New Jersey: Rowman & Littlefield Publishers, pp. 14–15.

References about efficient market theory:

Fama, E., 1970. "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance* 25, 383–417.

Fama, E., 1991. "Efficient Capital Markets: II." *Journal of Finance* 46, 1575–1617.

Jarrow, R. A., 1988. *Finance Theory*. Englewood Cliffs, NJ: Prentice-Hall.

QUESTIONS

Question 1

What are the four assumptions that characterize a lognormal distribution for stock price returns?

Question 2

The expected value of a continuously compounded rate of return is 12 percent per year and its volatility is 30 percent per year.

a) What is the expected return and volatility over a one-month period?
b) What is the expected return and volatility over a two-month period?
c) What is the expected return and volatility over a three-month period?
d) What is the expected return and volatility over a six-month period?

Question 3

What is the dollar expected return and standard deviation of the dollar return over $[t, t + \Delta]$ for the stock price in Expression (4.8)?

Question 4

Consider the following binomial model for a stock S_t, for $t = 0, 1, 2$.

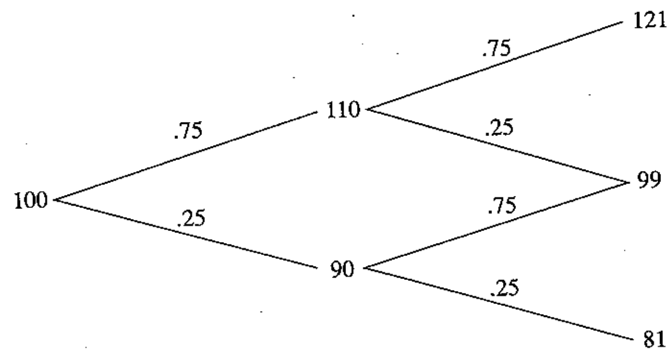

a) What is the probability that $S_2 = 121$?
b) What is the probability that $S_2 = 99$?
c) What is the probability that $S_2 = 81$?
d) What is the expected stock price at date 1?
e) What is the variance of the stock price at date 1? At date 2? Is the stock price variance increasing, decreasing, or constant across time?

Question 5

Consider a lognormal distribution with mean return per year of $\mu = 0.05$ and return standard deviation per year of $\sigma = 0.2$.

a) What are the up (U) and down (D) magnitudes for the binomial approximations to this lognormal distribution for an arbitrary set size Δ?
b) Compute the values in a) for $\Delta = 1, 1/2, 1/4, 1/8$. What happens to U and D as Δ decreases?
c) Using $\Delta = 1$, construct the binomial tree for two time steps. Let the initial stock price be \$100.

Question 6

The expected value of the continuously compounded rate of return is 12 percent per annum and its volatility is 30 percent per annum. The current stock price is \$100.

a) If the interval Δ is chosen to be one day, that is, $\Delta = 1/365$, use Expression (4.10) to compute the binomial distribution of stock prices over the interval.
b) Calculate the expected value of the stock price relative.
c) Calculate the volatility of the stock price relative. What relationship does your computed value have with the value of $\sigma\sqrt{\Delta}$?

Question 7

The expected value of the continuously compounded rate of return is 15 percent per annum and its volatility is 30 percent per annum. The current stock price is $100.

a) If the time interval is chosen to be one week, that is, $\Delta = 7/365$, use Expression (4.10) to compute a binomial lattice of stock prices over three weeks.
b) Compute the expected stock price at the end of one week.
c) Compute the expected stock price at the end of two weeks.
d) Compute the expected stock price at the end of three weeks.
e) How does your answer to part d) compare to the expected stock price computed using Expression (4.7)?

Question 8

The expected value of the continuously compounded rate of return is 18 percent per year and its volatility is 35 percent per year. The return is normally distributed. Below you are given the outcome of tossing a fair coin with outcomes Heads (H) or Tails (T).

Sequence	1	2	3	4	5	6	7	8	9	10
Outcome	T	T	H	T	H	H	T	H	H	H

The current stock price is $25.

a) If the interval is 5 days, that is, $\Delta = 5/365$, use Expression (4.10) to generate a random sequence of stock prices.
b) Repeat this exercise with the interval being one day, $\Delta = 1/365$.

Question 9

Given the information in Question 8, you are now given the values of a random drawing from a normal distribution with zero mean and variance Δ, where $\Delta = 1/365$.

Sequence	ΔW
1	−0.0604
2	−0.0219
3	0.0178
4	−0.0174
5	0.0244
6	0.0538
7	−0.0230
8	0.0152
9	0.0467
10	0.0396

a) Use Expression (4.11) to generate a random sequence of stock prices.
b) Compare your answers to those in Question 8b).

Question 10

Suppose the stock price is described via the stochastic differential equation

$$d\ln S(t) = \mu dt + \sigma dW(t).$$

a) What distribution does the stock price follow?
b) This description implicitly assumes that trading takes place continuously in time. What problems are there with this assumption?
c) This description implies that stock prices change continuously. What problems are there with this implication?

APPENDIX: THE EXPECTED VALUE OF THE FUTURE STOCK PRICE

We want to prove

$$E[S(T) \mid (S(0)] = S(0)\exp(\mu T + \sigma^2 T/2). \tag{A1}$$

PROOF From Expression (4.4) we have

$$S(T) = S(0)\exp[Z(T)],$$

where $Z(T)$ is normally distributed with mean μT, using Expression (4.5), and variance $\sigma^2 T$, using Expression (4.6). The probability density function for $Z(T)$ is

$$\frac{1}{\sigma_1\sqrt{2\pi}}\exp\left[-\frac{1}{2}\left(\frac{z-\mu_1}{\sigma_1}\right)^2\right],$$

where z is a realized value of $Z(T)$, $\mu_1 = \sigma T$, and $\sigma_1^2 = \sigma^2 T$.

The expected value of $S(T)$ conditional upon knowing $S(0)$ is

$$E[S(T) \mid S(0)] = S(0)\frac{1}{\sigma_1\sqrt{2\pi}}\int_{-\infty}^{+\infty}\exp(z)\exp\left[-\frac{1}{2}\left(\frac{z-\mu_1}{\sigma_1}\right)^2\right]dz. \tag{A2}$$

Completing the square gives

$$z - \frac{1}{2}\left(\frac{z-\mu_1}{\sigma_1}\right)^2 = \mu_1 + \sigma_1^2/2 - \frac{1}{2}\left[\frac{z-(\mu_1+\sigma_1^2)}{\sigma_1}\right]^2,$$

so that

$$E[S(T) \mid (S(0)] =$$
$$S(0)\exp(\mu_1 + \sigma_1^2/2)\frac{1}{\sigma_1\sqrt{2\pi}}\int_{-\infty}^{+\infty}\exp\left\{-\frac{1}{2}\left[\frac{z-(\mu_1 + \sigma_1^2)}{\sigma_1}\right]^2\right\}dz. \quad \text{(A3)}$$

Let $u \equiv [z - (\mu_1 + \sigma_1^2)]/\sigma_1$ so that

$$E[S(T) \mid (S(0)] = S(0)\exp(\mu_1 + \sigma_1^2/2)\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{+\infty}\exp\left(-\frac{1}{2}u^2\right)du. \quad \text{(A4)}$$

Now the last term on the right side is the area under the normal probability density function and equals unity:

$$\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{+\infty}\exp\left(-\frac{1}{2}u^2\right)du = 1. \quad \text{(A5)}$$

Therefore, after substituting for μ_1 and σ_1^2, we have the required result:

$$E[S(T) \mid (S(0)] = S(0)\exp(\mu T + \sigma^2 T/2). \quad \text{(A6)}$$

■