9 A Model of the Behavior of Stock Prices

Any variable whose value changes over time in an uncertain way is said to follow a *stochastic process*. Stochastic processes can be classified as "discrete time" or "continuous time." A discrete-time stochastic process is one where the value of the variable can only change at certain fixed points in time, whereas a continuous-time stochastic process is one where changes can take place at any time. Stochastic processes can also be classified as "continuous variable" or "discrete variable." In a continuous-variable process, the underlying variable can take any value within a certain range, whereas in a discrete-variable process, only certain discrete values are possible.

In this chapter we derive a continuous-variable, continuous-time stochastic process for stock prices. An understanding of this process is the first step to understanding the pricing of options and other more complicated derivative securities. It should be pointed out that in practice we do not observe stock prices following continuous-variable, continuous-time processes. Stock prices are restricted to discrete values (usually multiples of $\$\frac{1}{8}$) and changes can be observed only when the exchange is open. Nevertheless, the continuous-variable, continuous-time process proves to be a useful model for most purposes.

In this chapter and the next we introduce the reader to what is known as *stochastic calculus*. This is an extension of regular calculus that deals with continuous-time stochastic processes. Many people feel that continuous-time stochastic processes are so complicated that they must be left entirely to "rocket scientists." This is not so. The biggest hurdle to understanding these processes is the notation. In this chapter we present a step-by-step approach aimed at getting the reader over this hurdle.

9.1 THE MARKOV PROPERTY

A *Markov process* is a particular type of stochastic process where only the present value of a variable is relevant for predicting the future. The past history of the variable and the way in which the present has emerged from the past are irrelevant.

Stock prices are usually assumed to follow a Markov process. Suppose that the price of IBM stock is $100 now. If the stock price follows a Markov process, our predictions for the future should be unaffected by the price 1 week ago, 1 month ago, or 1 year ago. The only relevant piece of information is the fact that the price is now $100.[1] Predictions for the future are uncertain and must be expressed in terms of probability distributions. The Markov property implies that the probability distribution of the price at any particular future time depends only on the current stock price of $100.

The Markov property of stock prices is consistent with the weak form of market efficiency. This states that the present price of a stock impounds all the information contained in a record of past prices. If the weak form of market efficiency were not true, technical analysts could make above-average returns by interpreting charts of the past history of stock prices. There is very little evidence that they are in fact able to do this.

It is competition in the marketplace which tends to ensure that weak-form market efficiency holds. The very fact that there are many, many investors watching the stock market closely and trying to make a profit from it leads to a situation where a stock price at any given time impounds the information in past prices. Suppose that it is discovered that a particular pattern in past stock prices always gives a 65 percent chance of price rises in the near future. Investors would attempt to buy a stock as soon as the pattern was observed, and demand for the stock would immediately rise. This would lead to an immediate rise in its price and the observed effect would be eliminated—as would any profitable trading opportunities.

[1] Statistical properties of the stock price history of IBM may be useful in determining the characteristics of the stochastic process followed by the stock price (e.g., its volatility). The point which is being made here is that the particular path followed by the stock in the past is irrelevant.

9.2 WIENER PROCESSES

Models of stock price behavior are usually expressed in terms of what are known as *Wiener processes*. A Wiener process is a particular type of Markov stochastic process. It has been used in physics to describe the motion of a particle that is subject to a large number of small molecular shocks and is sometimes referred to as Brownian motion.

The behavior of a variable, z, which follows a Wiener process, can be understood by considering the changes in its value in small intervals of time. Consider a small interval of time of length Δt and define Δz as the change in z during Δt. There are two basic properties Δz must have for z to be following a Wiener process:

PROPERTY 1

Δz is related to Δt by the equation

$$\Delta z = \epsilon \sqrt{\Delta t} \tag{9.1}$$

where ϵ is a random drawing from a standardized normal distribution (i.e., a normal distribution with a mean of zero and a standard deviation of 1.0).

PROPERTY 2

The values of Δz for any two different short intervals of time Δt are independent.

It follows from the property 1 that Δz itself has a normal distribution with

$$\text{mean of } \Delta z = 0$$

$$\text{standard deviation of } \Delta z = \sqrt{\Delta t}$$

$$\text{variance of } \Delta z = \Delta t$$

Property 2 implies that z follows a Markov process.

Consider next the increase in the value of z during a relatively long period of time, T. We will denote this by $z(T) - z(0)$. It can be regarded as the sum of the increases in z in N small time intervals of length Δt, where

$$N = \frac{T}{\Delta t}$$

Thus

$$z(T) - z(0) = \sum_{i=1}^{N} \epsilon_i \sqrt{\Delta t} \tag{9.2}$$

where the ϵ_i $(i = 1, 2, \ldots, N)$ are random drawings from a standardized normal distribution. From property 2 the ϵ_i's are independent of each other. It follows from Equation (9.2) that $z(T) - z(0)$ is normally distributed with[2]

$$\text{mean of } [z(T) - z(0)] = 0$$

$$\text{variance of } [z(T) - z(0)] = N\,\Delta t = T$$

$$\text{standard deviation of } [z(T) - z(0)] = \sqrt{T}$$

Thus in any time interval of length T, the increase in the value of a variable that follows a Wiener process is normally distributed with a mean of zero and a standard deviation of $\sqrt{T}$. It should now be clear why Δz is defined as the product of ϵ and $\sqrt{\Delta t}$ rather than as the product of ϵ and Δt. Variances are additive for independent normal distributions; standard deviations are not. It makes sense to define the stochastic process so that the variance rather than the standard deviation of changes is proportional to the length of the time interval considered.

Example 9.1

Suppose that the value, z, of a variable which follows a Wiener process is initially 25 and that time is measured in years. At the end of 1 year the value of the variable is normally distributed with a mean of 25 and a standard deviation of 1.0. At the end of 2 years it is normally distributed with a mean of 25 and a standard deviation of $\sqrt{2}$, or 1.414. Note that our uncertainty about the value of the variable at a certain time in the future, as measured by its standard deviation, increases as the square root of how far we are looking ahead.

In ordinary calculus, it is usual to proceed from small changes to the limit as the small changes become closer to zero. Thus $\Delta y/\Delta x$ becomes dy/dx in the limit, and so on. We can proceed similarly when dealing with continuous-time stochastic processes. A Wiener process is the limit as $\Delta t \longrightarrow 0$ of the process described above for z. Figure 9.1 illustrates what happens to the path followed by z as the limit $\Delta t \longrightarrow 0$ is taken. Analogously to ordinary calculus, we write the limiting case of Equation (9.1) as

$$dz = \epsilon \sqrt{dt}$$

Generalized Wiener Process

The basic Wiener process that has been developed so far has a drift rate of zero and a variance rate of 1.0. The drift rate of zero means that the expected

[2]This result is based on the following well-known property of normal distributions. If a variable Y is equal to the sum of N independent normally distributed variables X_i $(1 \le i \le N)$, Y is itself normally distributed. The mean of Y is equal to the sum of the means of the X_i's. The variance of Y is equal to the sum of the variances of the X_i's.

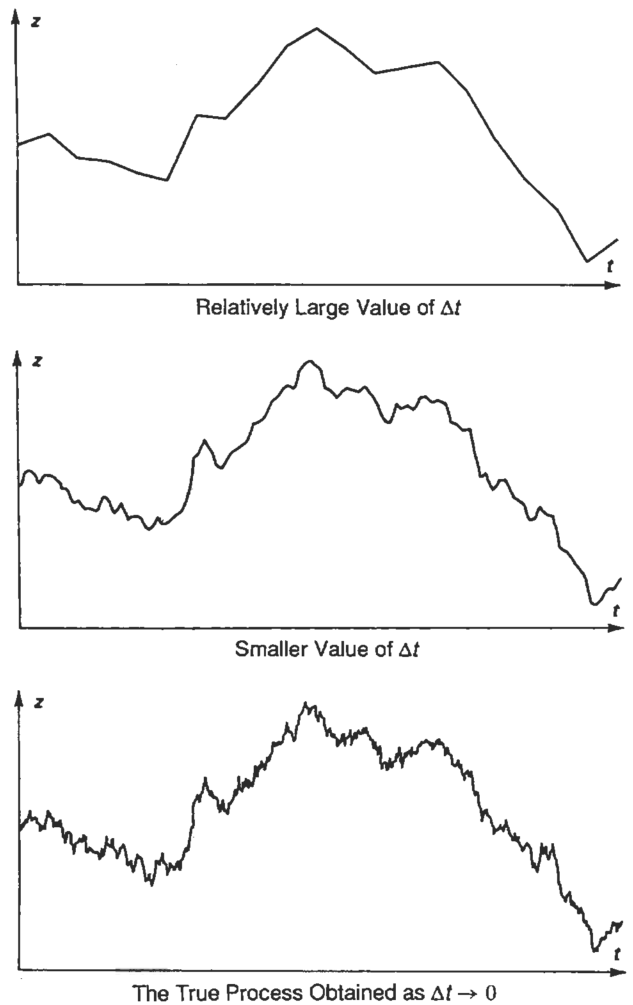

Figure 9.1 Illustration of How a Wiener Process is Obtained when $\Delta t \longrightarrow 0$ in Equation (9.1)

value of z at any future time is equal to its current value. The variance rate of 1.0 means that the variance of the change in z in a time interval of length T is $1.0 \times T$. A *generalized Wiener process* for a variable x can be defined in terms of dz as follows:

$$dx = a\,dt + b\,dz \tag{9.3}$$

where a and b are constants.

To understand Equation (9.3) it is useful to consider the two components on the right-hand side separately. The $a\,dt$ term implies that x has an expected drift rate of a per unit time. Without the $b\,dz$ term, the equation is

$$dx = a\,dt$$

which implies that

$$\frac{dx}{dt} = a$$

or

$$x = x_0 + at$$

where x_0 is the value of x at time zero. In a time interval of length T, x increases by an amount aT. The $b\,dz$ term on the right-hand side of Equation (9.3) can be regarded as adding noise or variability to the path followed by x. The amount of this noise or variability is b times a Wiener process. In a small time interval Δt, the change in the value of x, Δx, is from equations (9.1) and (9.3) given by:

$$\Delta x = a\,\Delta t + b\epsilon\sqrt{\Delta t}$$

where, as before, ϵ is a random drawing from a standardized normal distribution. Thus Δx has a normal distribution with

$$\text{mean of } \Delta x = a\,\Delta t$$

$$\text{standard deviation of } \Delta x = b\sqrt{\Delta t}$$

$$\text{variance of } \Delta x = b^2\,\Delta t$$

Similar arguments to those just given show that the change in the value of x in any time interval T is normally distributed with

$$\text{mean of change in } x = aT$$

$$\text{standard deviation of change in } x = b\sqrt{T}$$

$$\text{variance of change in } x = b^2T$$

Thus the generalized Wiener process given in Equation (9.3) has an expected drift rate (i.e., average drift per unit time) of a and a variance rate (i.e., variance per unit of time) of b^2. It is illustrated in Figure 9.2.

Example 9.2

Consider the situation where the cash position of a company, measured in thousands of dollars, follows a generalized Wiener process with a drift of 20 per year and a variance rate of 900 per year. Initially, the cash position is 50. At the end of 1 year the cash position will have a normal distribution with a mean of 70 and a standard deviation of $\sqrt{900}$ or 30. At the end of 6 months it will have a normal distribution with a mean of 60 and a standard

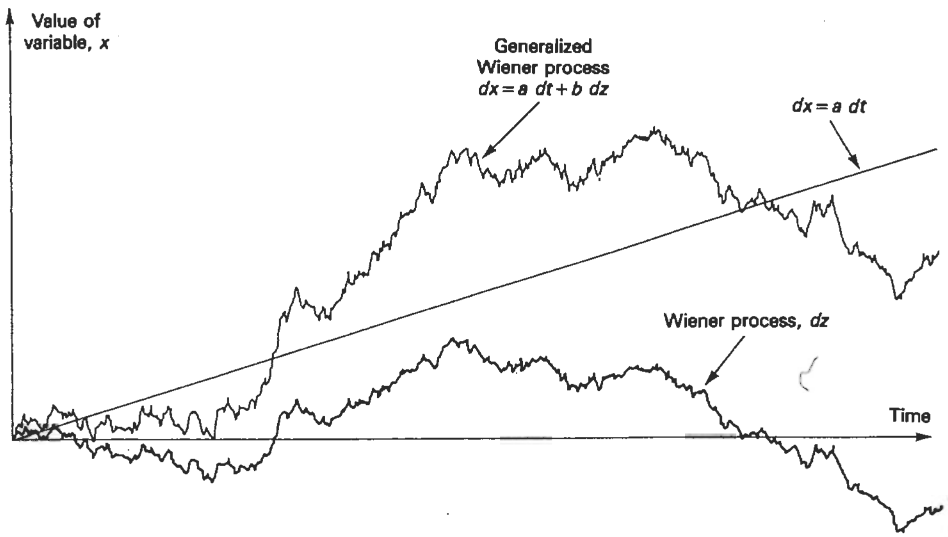

Figure 9.2 Generalized Wiener Process; a=0.3, b=1.5

deviation of $30\sqrt{0.5} = 21.21$. Note that our uncertainty about the cash position at some time in the future, as measured by its standard deviation, increases as the square root of how far ahead we are looking. Also, note that the cash position can become negative (which we can interpret as a situation where the company is borrowing funds).

Ito Process

A further type of stochastic process can be defined. This is known as an *Ito process*. It is a generalized Wiener process where the parameters a and b are functions of the value of the underlying variable, x, and time, t. Algebraically, an Ito process can be written

$$dx = a(x, t)dt + b(x, t)\,dz \tag{9.4}$$

Both the expected drift rate and variance rate of an Ito process are liable to change over time.

9.3 THE PROCESS FOR STOCK PRICES

In this section we discuss the stochastic process followed by the price of a non-dividend-paying stock. The effects of dividends on the process will be discussed in Chapter 10.

It is tempting to suggest that a stock price follows a generalized Wiener process; that is, that it has a constant expected drift rate and a constant variance rate. However, this model fails to capture a key aspect of stock prices. This is that the expected percentage return required by investors from a stock is independent of the stock's price. If investors require a 14 percent per annum expected return when the stock price is $10, then, *ceteris paribus*, they will also require a 14 percent per annum expected return when it is $50.

Clearly, the constant expected drift-rate assumption is inappropriate and needs to be replaced by the assumption that the expected drift, expressed as a proportion of the stock price, is constant. The latter implies that if S is the stock price, the expected drift rate in S is μS for some constant parameter, μ. Thus, in a short interval of time, Δt, the expected increase in S is $\mu S\,\Delta t$. The parameter, μ, is the expected rate of return on the stock, expressed in decimal form.

If the variance rate of the stock price is always zero, this model implies that

$$dS = \mu S\,dt$$

or

$$\frac{dS}{S} = \mu\,dt$$

so that

$$S = S_0 e^{\mu t} \tag{9.5}$$

where S_0 is the stock price at time zero. Equation (9.5) shows that, when the variance rate is zero, the stock price grows at a continuously compounded rate of μ per unit time.

In practice, of course, a stock price does exhibit volatility. A reasonable assumption is that the variance of the percentage return in a short period of time, Δt, is the same regardless of the stock price. In other words, an investor is just as uncertain as to his or her percentage return when the stock price is $50 as when it is $10. Define σ^2 as the variance rate of the proportional change in the stock price. This means that $\sigma^2\,\Delta t$ is the variance of the proportional change in the stock price in time Δt and that $\sigma^2 S^2\,\Delta t$ is the variance of the actual change in the stock price, S, during Δt. The instantaneous variance rate of S is therefore $\sigma^2 S^2$.

These arguments suggest that S can be represented by an Ito process which has instantaneous expected drift rate μS and instantaneous variance rate $\sigma^2 S^2$. This can be written as

$$dS = \mu S\,dt + \sigma S\,dz$$

or:

$$\frac{dS}{S} = \mu\,dt + \sigma\,dz \tag{9.6}$$

Equation (9.6) is the most widely used model of stock price behavior. The variable σ is usually referred to as the *stock price volatility*. The variable μ is its expected rate of return.

Example 9.3

Consider a stock that pays no dividends, has a volatility of 30% per annum, and provides an expected return of 15% per annum with continuous compounding. In this case $\mu = 0.15$ and $\sigma = 0.30$. The process for the stock price is

$$\frac{dS}{S} = 0.15\,dt + 0.30\,dz$$

If S is the stock price at a particular time and ΔS is the increase in the stock price in the next small interval of time,

$$\frac{\Delta S}{S} = 0.15\,\Delta t + 0.30\epsilon\sqrt{\Delta t}$$

where ϵ is a random drawing from a standardized normal distribution. Consider a time interval of 1 week or 0.0192 year and suppose that the initial stock price is \$100. Then $\Delta t = 0.0192$, $S = 100$, and

$$\Delta S = 100(0.00288 + 0.0416\epsilon)$$

showing that the price increase is a random drawing from a normal distribution with mean \$0.288 and standard deviation \$4.16.

9.4 A REVIEW OF THE MODEL

The model of stock price behavior that has been developed in this chapter [see Equation (9.6)] is sometimes known as *geometric Brownian motion*. The discrete-time version of the model is

$$\frac{\Delta S}{S} = \mu\,\Delta t + \sigma\epsilon\sqrt{\Delta t} \tag{9.7}$$

The variable ΔS is the change in the stock price, S, in a small interval of time, Δt; and ϵ is a random drawing from a standardized normal distribution (i.e., a normal distribution with a mean of zero and standard deviation of 1.0). The parameter μ is the expected rate of return per unit time from the stock and the parameter σ is the volatility of the stock price. Both of these parameters are assumed constant.

The left-hand side of Equation (9.7) is the proportional return provided by the stock in a short period of time Δt. The term $\mu\,\Delta t$ is the expected value of this return, while the term $\sigma\epsilon\sqrt{\Delta t}$ is the stochastic component of the return. The variance of the stochastic component (and therefore of the whole return) is $\sigma^2\,\Delta t$.

Equation (9.7) shows that $\Delta S/S$ is normally distributed with mean $\mu\,\Delta t$ and standard deviation $\sigma\sqrt{\Delta t}$. In other words,

$$\frac{\Delta S}{S} \sim \phi(\mu\,\Delta t,\ \sigma\sqrt{\Delta t}) \tag{9.8}$$

where $\phi(m, s)$ denotes a normal distribution with mean m and standard deviation s.

MONTE CARLO SIMULATION

Suppose that the expected return from the stock is 14 percent per annum and that the standard deviation of the return (i.e., the volatility) is 20 percent per annum. If time is measured in years, it follows that

$$\mu = 0.14$$

$$\sigma = 0.20$$

Suppose that $\Delta t = 0.01$ so that we are considering changes in the stock price in time intervals of length 0.01 year (or 3.65 days). It follows that $\Delta S/S$ is normal with mean 0.0014 (= 0.14×0.01) and standard deviation 0.02 (= $0.2 \times \sqrt{0.01}$), that is,

$$\frac{\Delta S}{S} \sim \phi(0.0014,\ 0.02) \tag{9.9}$$

A path for the stock price can be simulated by sampling repeatedly from $\phi(0.0014, 0.02)$. One procedure for doing this is to sample values, v_1, from a standardized normal distribution [i.e., $\phi(0, 1)$] and then convert these to samples, v_2, from $\phi(0.0014, 0.02)$ using

$$v_2 = 0.0014 + 0.02v_1 \tag{9.10}$$

Table 9.1 shows one particular simulation of stock price movements. The initial stock price is assumed to be \$20. For the first period the random number, v_1, sampled from $\phi(0, 1)$ is 0.52. Using Equation (9.10) this gives a random sample of 0.0118 from $\phi(0.0014, 0.02)$. Using Equation (9.9), $\Delta S = 20 \times 0.0118$, or 0.236. At the beginning of the next period the stock price is therefore \$20.236, and so on. Note that the samples, v_2, must be independent of each other. Otherwise, the Markov property, discussed in Section 9.1, does not hold.

Table 9.1 assumes that stock prices are measured to the nearest 0.001, which of course is not the case. To get the stock price that would be quoted, the figures in the first column of the table should be rounded to the nearest $\$\frac{1}{8}$. It is important to realize that the table shows only one possible pattern of stock price movements. Different random samples would lead to different price movements. Any small time interval Δt can be used in the simulation. However, only in the limit as $\Delta t \longrightarrow 0$ is a true description of geometric Brownian motion obtained. The final stock price

TABLE 9.1 Simulation of Stock Price when $\mu = 0.14$ and $\sigma = 0.20$ During Periods of Length 0.01 Year

Stock Price at Start of Period	Random Sample, v_1, from $\phi(0, 1)$	Corresponding Random Sample, v_2, from $\phi(0.0014, 0.02)$	Change in Stock Price During Period
20.000	0.52	0.0118	0.236
20.236	1.44	0.0302	0.611
20.847	−0.86	−0.0158	−0.329
20.518	1.46	0.0306	0.628
21.146	−0.69	−0.0124	−0.262
20.883	−0.74	−0.0134	−0.280
20.603	0.21	0.0056	0.115
20.719	−1.10	−0.0206	−0.427
20.292	0.73	0.0160	0.325
20.617	1.16	0.0246	0.507
21.124	2.56	0.0526	1.111

of 21.124 in Table 9.1 can be regarded as a random sample from the distribution of stock prices at the end of 10 time intervals or one-tenth of a year. By repeatedly simulating movements in the stock price, as in Table 9.1, a complete probability distribution of the stock price at the end of one-tenth of a year is obtained.

9.5 THE PARAMETERS

The process for stock prices that has been developed in this chapter involves two parameters, μ and σ. The values of those parameters depend on the units in which time is measured. Here and elsewhere in this book, we assume that time is measured in years.

The parameter μ is the expected proportional return earned by an investor in a short period of time. It is annualized and expressed as a proportion. Most investors require higher expected returns to induce them to take higher risks. It follows that the value of μ should depend on the risk of the return from the stock.[3] It should also depend on the level of interest rates in the economy. The higher the level of interest rates, the higher the expected return required on any given stock. On average, μ is about 8 percent greater than the return on a risk-free investment such as a Treasury bill.[4] Thus when the return on Treasury bills is 8 percent per annum, or 0.08, a typical value of μ is 0.16; that is, a typical expected return on a stock is 16 percent per annum.

[3]More precisely, μ depends on that part of the risk which cannot be diversified away by the investor.

[4]See R. G. Ibbotson and R. A. Sinquefield, *Stocks, Bonds, Bills and Inflation: The Past and the Future* (Charlotteville, Va.: Financial Analyst Research Foundation, 1982), Exhibit 29, p. 71.

Fortunately, we do not have to concern ourselves with the determinants of μ in any detail because the value of a derivative security dependent on a stock is in general independent of μ. The parameter σ, the stock price volatility, is, by contrast, critically important to the determination of the value of most contingent claims. Procedures for estimating σ empirically are discussed in Chapter 10. Typical values of σ for a stock are in the range 0.20 to 0.40 (i.e., 20 percent to 40 percent).

The standard deviation of the proportional change in the stock price in a small interval of time Δt is $\sigma\sqrt{\Delta t}$. As a rough approximation, the standard deviation of the proportional change in the stock price in a relatively long period of time, T, is $\sigma\sqrt{T}$. This means that, as an approximation, volatility can be interpreted as the standard deviation of the change in the stock price in one year.

Note that the standard deviation of the proportional change in the stock price in a relatively long time interval, T, is not exactly $\sigma\sqrt{T}$. This is because proportional changes are not additive. (For example, a 10 percent increase in a stock price followed by a 20 percent increase leads to a total increase of 32 percent, not 30 percent.) In Chapter 10, the probability distribution of the change in the stock price in a relatively long time period T will be shown to be lognormal. Also, the volatility of a stock price will be shown to be exactly equal to the standard deviation of the continuously compounded return provided by the stock in one year.

9.6 A BINOMIAL MODEL

At various points in this book we will use a binomial model as a discrete-time representation of the continuous-time model for stock prices which has been described in this chapter. Suppose that the stock price starts at S. Under the binomial model, the stock price follows the process illustrated in Figure 9.3 in the next small time interval of length Δt. It moves up to Su with probability p and down to Sd with probability $1 - p$. Figure 9.4 illustrates how the binomial model leads to

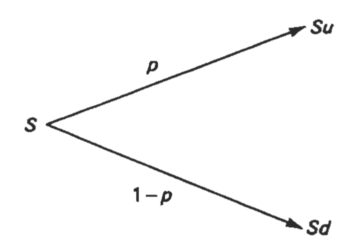

Figure 9.3 Binomial Model

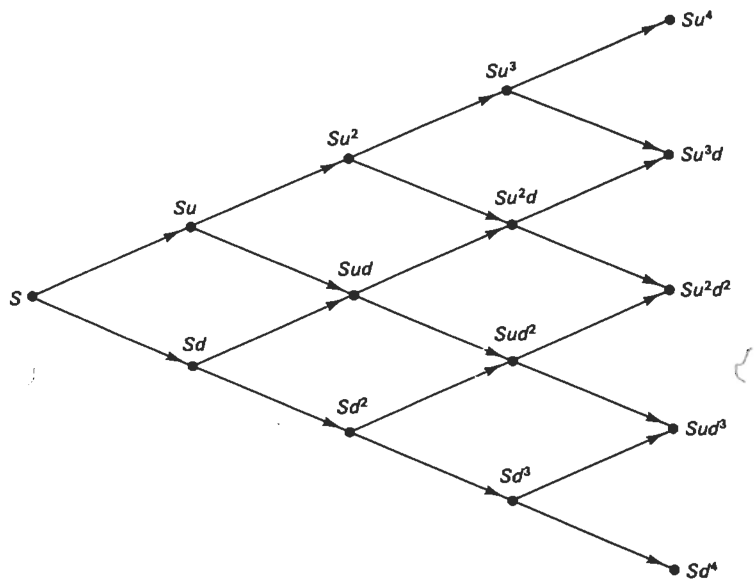

Figure 9.4 Stock Price Movements over Four Time Periods Using Binomial Model

three alternative stock prices at the end of two time intervals, four alternative stock prices at the end of three time intervals, and so on.

The variables u, d, and p must be chosen so that, for a small Δt, the expected return from the stock price in time Δt is $\mu\,\Delta t$ and the variance of the return in time Δt is $\sigma^2\,\Delta t$. One way of doing this is by setting[5]

$$u = e^{\sigma\sqrt{\Delta t}}, \qquad d = \frac{1}{u}, \qquad p = \frac{e^{\mu\Delta t} - d}{u - d}$$

[5]To demonstrate that these values of u, d, and p have the right properties, note that the expected stock price at time Δt is

$$pSu + (1-p)Sd = Se^{\mu\Delta t}$$

The variance of the stock price at time Δt is

$$pS^2u^2 + (1-p)S^2d^2 - (Se^{\mu\Delta t})^2$$

This equals

$$S^2\left[e^{\mu\Delta t}(e^{\sigma\sqrt{\Delta t}} + e^{-\sigma\sqrt{\Delta t}}) - 1 - e^{2\mu\Delta t}\right]$$

Expanding e^x in series form $e^x = 1 + x + x^2/2 + x^3/6 + \ldots$, the variance of the stock price is $S^2\sigma^2\,\Delta t$ when terms of order Δt^2 and higher are ignored.

It can be shown that in the limit as $\Delta t \longrightarrow 0$, this binomial model of stock price movements becomes the geometric Brownian motion model which has been developed in this chapter.

Example 9.4

Consider a stock price that provides an expected return of 12% per annum and has a volatility of 30% per annum. Suppose that the binomial model is used to represent movements in time periods of 0.04 year (approximately 2 weeks). In this case $\mu = 0.12$, $\sigma = 0.30$, and $\Delta t = 0.04$ and, from the previous equations:

$$u = e^{0.30\times\sqrt{0.04}} = 1.0618$$

$$d = \frac{1}{u} = 0.9418$$

$$p = \frac{e^{0.12\times 0.04} - 0.9418}{1.0618 - 0.9418} = 0.525$$

If the stock price starts at \$100, possible movements over four time intervals of length Δt are as illustrated in Figure 9.5. The probability of an up movement is always 0.525 and the probability of a down movement is always 0.475. For the stock price of \$112.7 to occur at the end of the four time intervals there must be three up movements and one down movement. There are four ways that this can happen. These are $DUUU$, $UDUU$, $UUDU$, and $UUUD$, where U denotes an up movement and D denotes a down movement. Hence

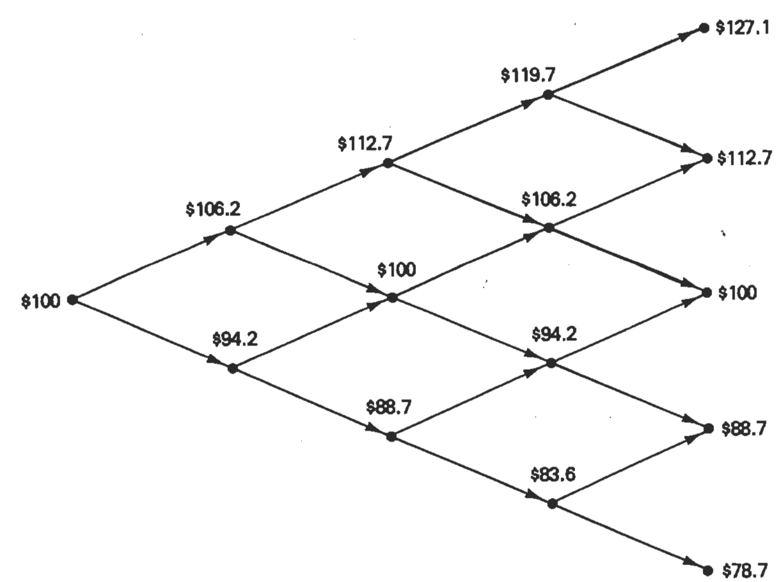

Figure 9.5 Stock Price Movements in Example 9.4

the probability of the stock price being $112.7 at the end of four time intervals is

$$4 \times 0.525^3 \times 0.475 = 0.275$$

The probabilities of stock prices of $127.1, $100.0, $88.7, and $78.7 can similarly be shown to be 0.076, 0.373, 0.225, and 0.051, respectively.

9.7 SUMMARY

Stochastic processes describe the probabilistic evolution of the value of a variable through time. A Markov process is one where only the present value of the variable is relevant for predicting the future. The past history of the variable and the way in which the present has emerged from the past is irrelevant.

A Wiener process, dz, is a process describing the evolution of a normally distributed variable. The drift of the process is zero and the variance rate is 1 per unit time. This means that, if the value of the variable is x at time zero, at time T it is normally distributed with mean x and standard deviation $\sqrt{T}$.

A generalized Wiener process describes the evolution of a normally distributed variable with a drift of a per unit time and a variance rate of b^2 per unit time where a and b are constants. This means that if the value of the variable is x at time zero, at time T it is normally distributed with a mean of $x + aT$ and a standard deviation of $b\sqrt{T}$.

An Ito process is a process where the drift and variance rate of x can be a function of both x itself and time. The change in x in a very short period of time is normally distributed but its change over longer periods of time is liable to be non-normal.

In this chapter we have developed a plausible Markov stochastic process for the behavior of a stock price over time. The process is widely used in the valuation of derivative securities. It is known as geometric Brownian motion. Under this process, the proportional rate of return to the holder of the stock in any small interval of time is normally distributed and the returns in any two different small intervals of time are independent.

One way of gaining an intuitive understanding of a stochastic process for a variable is to simulate the behavior of the variable. This involves dividing a time interval into many small time steps and randomly sampling possible paths for the variable. The future probability distribution for the variable can then be calculated. Monte Carlo simulation will be discussed further in Chapter 14.

SUGGESTIONS FOR FURTHER READING

On the Markov Property of Stock Prices

BREALEY, R. A., *An Introduction to Risk and Return from Common Stock* (2nd ed.). Cambridge, Mass.: MIT Press, 1983.

COOTNER, P. H., ed., *The Random Character of Stock Market Prices*. Cambridge, Mass.: MIT Press, 1964.

On Stochastic Processes

COX, D. R., and H. D. MILLER, *The Theory of Stochastic Processes*. London: Chapman & Hall, 1965.

FELLER, W., *Probability Theory and its Applications*, vols. 1 and 2. New York: John Wiley, 1950.

QUESTIONS AND PROBLEMS

9.1. What would it mean to assert that the temperature at a certain place follows a Markov process? Do you think that temperatures do, in fact, follow a Markov process?

9.2. Can a trading rule based on the past history of a stock's price ever produce returns which are consistently above average? Discuss.

***9.3.** A company's cash position, measured in millions of dollars, follows a generalized Wiener process with a drift rate of 0.1 per month and a variance rate of 0.16 per month. The initial cash position is 2.0.
- **(a)** What are the probability distributions of the cash position after 1 month, 6 months, and 1 year?
- **(b)** What are the probabilities of a negative cash position at the end of 6 months and 1 year?
- **(c)** At what time in the future is the probability of a negative cash position greatest?

***9.4.** A company's cash position, measured in millions of dollars, follows a generalized Wiener process with a drift rate of 1.5 per quarter and a variance rate of 4.0 per quarter. How high does the company's initial cash position have to be for the company to have a less than 5% chance of a negative cash position by the end of 1 year?

***9.5.** Variables X_1 and X_2 follow generalized Wiener processes with drift rates μ_1 and μ_2 and variances σ_1^2 and σ_2^2. What process does $X_1 + X_2$ follow if:
- **(a)** The changes in X_1 and X_2 in any short interval of time are uncorrelated?
- **(b)** There is a correlation ρ between the changes in X_1 and X_2 in any short interval of time?

9.6. Consider a variable, S, which follows the process

$$dS = \mu\, dt + \sigma\, dz$$

For the first 3 years, $\mu = 2$ and $\sigma = 3$; for the next 3 years, $\mu = 3$ and $\sigma = 4$. If the initial value of the variable is 5, what is the probability distribution of the value of the variable at the end of year 6?

9.7. Suppose that a stock price has an expected return of 16% per annum and a volatility of 30% per annum. When the stock price at the end of a certain day is $50, calculate the following:
- **(a)** The expected stock price at the end of the next day.
- **(b)** The standard deviation of the stock price at the end of the next day.