

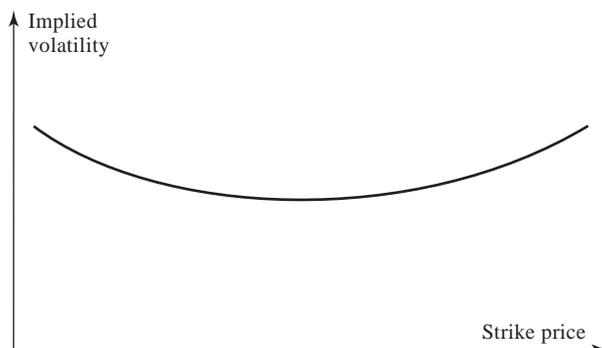
19.2 FOREIGN CURRENCY OPTIONS

The volatility smile used by traders to price foreign currency options has the general form shown in Figure 19.1. The implied volatility is relatively low for at-the-money options. It becomes progressively higher as an option moves either into the money or out of the money.

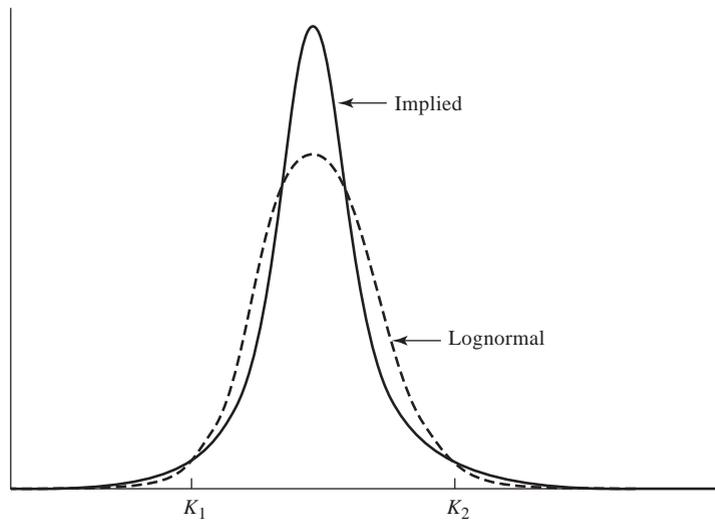
In the appendix at the end of this chapter, we show how to determine the risk-neutral probability distribution for an asset price at a future time from the volatility smile given by options maturing at that time. We refer to this as the *implied distribution*. The volatility smile in Figure 19.1 corresponds to the implied distribution shown by the solid line in Figure 19.2. A lognormal distribution with the same mean and standard deviation as the implied distribution is shown by the dashed line in Figure 19.2. It can be seen that the implied distribution has heavier tails than the lognormal distribution.¹

To see that Figures 19.1 and 19.2 are consistent with each other, consider first a deep-out-of-the-money call option with a high strike price of K_2 . This option pays off only if the exchange rate proves to be above K_2 . Figure 19.2 shows that the probability of this is higher for the implied probability distribution than for the lognormal distribution. We therefore expect the implied distribution to give a relatively high price for the option. A relatively high price leads to a relatively high implied volatility—and this is exactly what we observe in Figure 19.1 for the option. The two figures are therefore consistent with each other for high strike prices. Consider next a deep-out-of-the-money put option with a low strike price of K_1 . This option pays off only if the exchange rate proves to be below K_1 . Figure 19.2 shows that the probability of this is also higher for the implied probability distribution than for the lognormal distribution. We therefore expect the implied distribution to give a relatively high price, and a relatively high implied volatility, for this option as well. Again, this is exactly what we observe in Figure 19.1.

Figure 19.1 Volatility smile for foreign currency options.



¹ This is known as *kurtosis*. Note that, in addition to having a heavier tail, the implied distribution is more “peaked.” Both small and large movements in the exchange rate are more likely than with the lognormal distribution. Intermediate movements are less likely.

Figure 19.2 Implied and lognormal distribution for foreign currency options.

Empirical Results

We have just shown that the volatility smile used by traders for foreign currency options implies that they consider that the lognormal distribution understates the probability of extreme movements in exchange rates. To test whether they are right, Table 19.1 examines the daily movements in 12 different exchange rates over a 10-year period.² The first step in the production of the table is to calculate the standard deviation of daily percentage change in each exchange rate. The next stage is to note how often the actual percentage change exceeded 1 standard deviation, 2 standard deviations, and so on. The final stage is to calculate how often this would have happened if the percentage changes had been normally distributed. (The lognormal model implies that percentage changes are almost exactly normally distributed over a one-day time period.)

Table 19.1 Percentage of days when daily exchange rate moves are greater than 1, 2, . . . , 6 standard deviations (SD = standard deviation of daily change).

	<i>Real world</i>	<i>Lognormal model</i>
>1 SD	25.04	31.73
>2 SD	5.27	4.55
>3 SD	1.34	0.27
>4 SD	0.29	0.01
>5 SD	0.08	0.00
>6 SD	0.03	0.00

² The results in this table are taken from J.C. Hull and A. White, "Value at Risk When Daily Changes in Market Variables Are Not Normally Distributed." *Journal of Derivatives*, 5, No. 3 (Spring 1998): 9–19.

Business Snapshot 19.1 Making Money from Foreign Currency Options

Suppose that most market participants think that exchange rates are lognormally distributed. They will be comfortable using the same volatility to value all options on a particular exchange rate. You have just done the analysis in Table 19.1 and know that the lognormal assumption is not a good one for exchange rates. What should you do?

The answer is that you should buy deep-out-the-money call and put options on a variety of different currencies and wait. These options will be relatively inexpensive and more of them will close in the money than the lognormal model predicts. The present value of your payoffs will on average be much greater than the cost of the options.

In the mid-1980s a few traders knew about the heavy tails of foreign exchange probability distributions. Everyone else thought that the lognormal assumption of Black–Scholes–Merton was reasonable. The few traders who were well informed followed the strategy we have described—and made lots of money. By the late 1980s everyone realized that foreign currency options should be priced with a volatility smile and the trading opportunity disappeared.

Daily changes exceed 3 standard deviations on 1.34% of days. The lognormal model predicts that this should happen on only 0.27% of days. Daily changes exceed 4, 5, and 6 standard deviations on 0.29%, 0.08%, and 0.03% of days, respectively. The lognormal model predicts that we should hardly ever observe this happening. The table therefore provides evidence to support the existence of heavy tails (Figure 19.2) and the volatility smile used by traders (Figure 19.1). Business Snapshot 19.1 shows how you could have made money if you had done the analysis in Table 19.1 ahead of the rest of the market.

Reasons for the Smile in Foreign Currency Options

Why are exchange rates not lognormally distributed? Two of the conditions for an asset price to have a lognormal distribution are:

1. The volatility of the asset is constant.
2. The price of the asset changes smoothly with no jumps.

In practice, neither of these conditions is satisfied for an exchange rate. The volatility of an exchange rate is far from constant, and exchange rates frequently exhibit jumps.³ It turns out that the effect of both a nonconstant volatility and jumps is that extreme outcomes become more likely.

The impact of jumps and nonconstant volatility depends on the option maturity. As the maturity of the option is increased, the percentage impact of a nonconstant volatility on prices becomes more pronounced, but its percentage impact on implied volatility usually becomes less pronounced. The percentage impact of jumps on both prices and the implied volatility becomes less pronounced as the maturity of the option is increased.⁴ The result of all this is that the volatility smile becomes less pronounced as option maturity increases.

³ Sometimes the jumps are in response to the actions of central banks.

⁴ When we look at sufficiently long-dated options, jumps tend to get “averaged out,” so that the exchange rate distribution when there are jumps is almost indistinguishable from the one obtained when the exchange rate changes smoothly.

19.3 EQUITY OPTIONS

The volatility smile for equity options has been studied by Rubinstein (1985, 1994) and Jackwerth and Rubinstein (1996). Prior to 1987 there was no marked volatility smile. Since 1987 the volatility smile used by traders to price equity options (both on individual stocks and on stock indices) has had the general form shown in Figure 19.3. This is sometimes referred to as a *volatility skew*. The volatility decreases as the strike price increases. The volatility used to price a low-strike-price option (i.e., a deep-out-of-the-money put or a deep-in-the-money call) is significantly higher than that used to price a high-strike-price option (i.e., a deep-in-the-money put or a deep-out-of-the-money call).

The volatility smile for equity options corresponds to the implied probability distribution given by the solid line in Figure 19.4. A lognormal distribution with the same mean and standard deviation as the implied distribution is shown by the dotted line. It can be seen that the implied distribution has a heavier left tail and a less heavy right tail than the lognormal distribution.

To see that Figures 19.3 and 19.4 are consistent with each other, we proceed as for Figures 19.1 and 19.2 and consider options that are deep out of the money. From Figure 19.4, a deep-out-of-the-money call with a strike price of K_2 has a lower price when the implied distribution is used than when the lognormal distribution is used. This is because the option pays off only if the stock price proves to be above K_2 , and the probability of this is lower for the implied probability distribution than for the lognormal distribution. Therefore, we expect the implied distribution to give a relatively low price for the option. A relatively low price leads to a relatively low implied volatility—and this is exactly what we observe in Figure 19.3 for the option. Consider next a deep-out-of-the-money put option with a strike price of K_1 . This option pays off only if the stock price proves to be below K_1 . Figure 19.4 shows that the probability of this is higher for the

Figure 19.3 Volatility smile for equities.

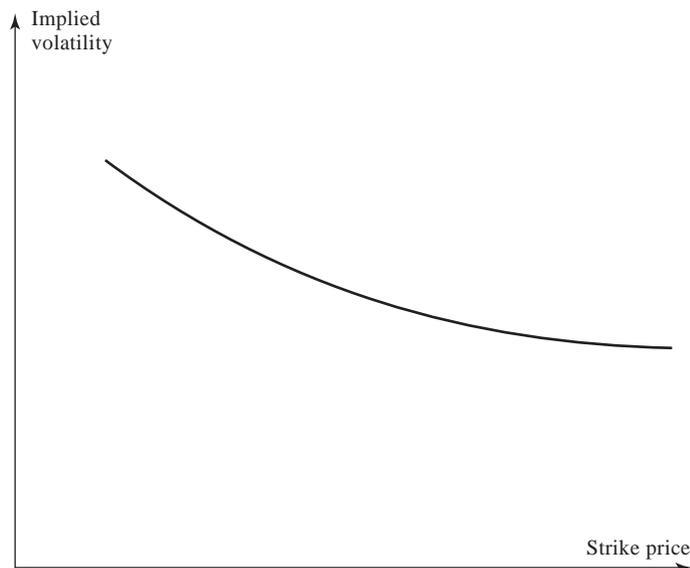
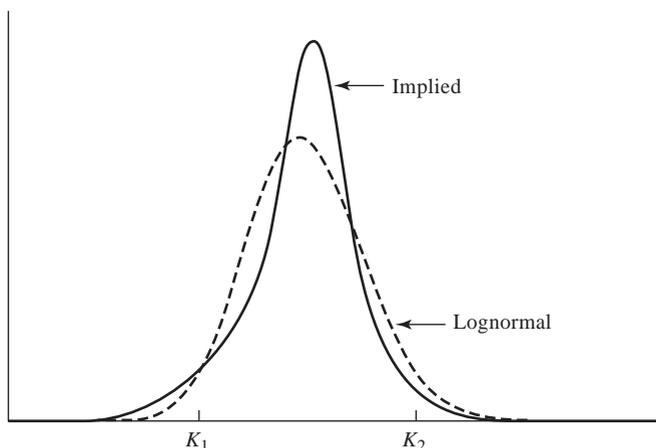


Figure 19.4 Implied distribution and lognormal distribution for equity options.



implied probability distribution than for the lognormal distribution. We therefore expect the implied distribution to give a relatively high price, and a relatively high implied volatility, for this option. Again, this is exactly what we observe in Figure 19.3.

The Reason for the Smile in Equity Options

One possible explanation for the smile in equity options concerns leverage. As a company's equity declines in value, the company's leverage increases. This means that the equity becomes more risky and its volatility increases. As a company's equity increases in value, leverage decreases. The equity then becomes less risky and its volatility decreases. This argument suggests that we can expect the volatility of a stock to be a decreasing function of the stock price and is consistent with Figures 19.3 and 19.4. Another explanation is "crashophobia" (see Business Snapshot 19.2).

19.4 ALTERNATIVE WAYS OF CHARACTERIZING THE VOLATILITY SMILE

So far we have defined the volatility smile as the relationship between implied volatility and strike price. The relationship depends on the current price of the asset. For example, the lowest point of the volatility smile in Figure 19.1 is usually close to the current exchange rate. If the exchange rate increases, the volatility smile tends to move to the right; if the exchange rate decreases, the volatility smile tends to move to the left. Similarly, in Figure 19.3, when the equity price increases, the volatility skew tends to move to the right, and when the equity price decreases, it tends to move to the left.⁵ For this reason the volatility smile is often calculated as the relationship between the implied volatility and K/S_0 rather than as the relationship between the implied volatility and K . The smile is then much more stable.

⁵ Research by Derman suggests that this adjustment is sometimes "sticky" in the case of exchange-traded options. See E. Derman, "Regimes of Volatility," *Risk*, April 1999: 55–59.

Business Snapshot 19.2 Crashophobia

It is interesting that the pattern in Figure 19.3 for equities has existed only since the stock market crash of October 1987. Prior to October 1987, implied volatilities were much less dependent on strike price. This has led Mark Rubinstein to suggest that one reason for the equity volatility smile may be “crashophobia.” Traders are concerned about the possibility of another crash similar to October 1987, and they price options accordingly.

There is some empirical support for this explanation. Declines in the S&P 500 tend to be accompanied by a steepening of the volatility skew. When the S&P increases, the skew tends to become less steep.

A refinement of this is to calculate the volatility smile as the relationship between the implied volatility and K/F_0 , where F_0 is the forward price of the asset for a contract maturing at the same time as the options that are considered. Traders also often define an “at-the-money” option as an option where $K = F_0$, not as an option where $K = S_0$. The argument for this is that F_0 , not S_0 , is the expected stock price on the option’s maturity date in a risk-neutral world.⁶

Yet another approach to defining the volatility smile is as the relationship between the implied volatility and the delta of the option (where delta is defined as in Chapter 18). This approach sometimes makes it possible to apply volatility smiles to options other than European and American calls and puts. When the approach is used, an at-the-money option is then defined as a call option with a delta of 0.5 or a put option with a delta of -0.5 . These are referred to as “50-delta options.”

19.5 THE VOLATILITY TERM STRUCTURE AND VOLATILITY SURFACES

Traders allow the implied volatility to depend on time to maturity as well as strike price. Implied volatility tends to be an increasing function of maturity when short-dated volatilities are historically low. This is because there is then an expectation that volatilities will increase. Similarly, volatility tends to be a decreasing function of maturity when short-dated volatilities are historically high. This is because there is then an expectation that volatilities will decrease.

Volatility surfaces combine volatility smiles with the volatility term structure to tabulate the volatilities appropriate for pricing an option with any strike price and any maturity. An example of a volatility surface that might be used for foreign currency options is given in Table 19.2.

One dimension of Table 19.2 is K/S_0 ; the other is time to maturity. The main body of the table shows implied volatilities calculated from the Black–Scholes–Merton model. At any given time, some of the entries in the table are likely to correspond to options for which reliable market data are available. The implied volatilities for these options are calculated directly from their market prices and entered into the table. The rest of the table is typically determined using interpolation. The table shows that the volatility smile becomes less pronounced as the option maturity increases. As mentioned earlier, this is

⁶ As explained in Chapter 27, whether the futures or forward price of the asset is the expected price in a risk-neutral world depends on exactly how the risk-neutral world is defined.

Table 19.2 Volatility surface.

	K/S_0				
	0.90	0.95	1.00	1.05	1.10
1 month	14.2	13.0	12.0	13.1	14.5
3 month	14.0	13.0	12.0	13.1	14.2
6 month	14.1	13.3	12.5	13.4	14.3
1 year	14.7	14.0	13.5	14.0	14.8
2 year	15.0	14.4	14.0	14.5	15.1
5 year	14.8	14.6	14.4	14.7	15.0

what is observed for currency options. (It is also what is observed for options on most other assets.)

When a new option has to be valued, financial engineers look up the appropriate volatility in the table. For example, when valuing a 9-month option with a K/S_0 ratio of 1.05, a financial engineer would interpolate between 13.4 and 14.0 in Table 19.2 to obtain a volatility of 13.7%. This is the volatility that would be used in the Black–Scholes–Merton formula or a binomial tree. When valuing a 1.5-year option with a K/S_0 ratio of 0.925, a two-dimensional (bilinear) interpolation would be used to give an implied volatility of 14.525%.

The shape of the volatility smile depends on the option maturity. As illustrated in Table 19.2, the smile tends to become less pronounced as the option maturity increases. Define T as the time to maturity and F_0 as the forward price of the asset for a contract maturing at the same time as the option. Some financial engineers choose to define the volatility smile as the relationship between implied volatility and

$$\frac{1}{\sqrt{T}} \ln\left(\frac{K}{F_0}\right)$$

rather than as the relationship between the implied volatility and K . The smile is then usually much less dependent on the time to maturity.⁷

19.6 GREEK LETTERS

The volatility smile complicates the calculation of Greek letters. Assume that the relationship between the implied volatility and K/S for an option with a certain time to maturity remains the same.⁸ As the price of the underlying asset changes, the implied volatility of the option changes to reflect the option's "moneyness" (i.e., the extent to which it is in or out of the money). The formulas for Greek letters given in Chapter 18

⁷ For a discussion of this approach, see S. Natenberg *Option Pricing and Volatility: Advanced Trading Strategies and Techniques*, 2nd edn. McGraw-Hill, 1994; R. Tompkins *Options Analysis: A State of the Art Guide to Options Pricing*, Burr Ridge, IL: Irwin, 1994.

⁸ It is interesting that this natural model is internally consistent only when the volatility smile is flat for all maturities. See, for example, T. Daghish, J. Hull, and W. Suo, "Volatility Surfaces: Theory, Rules of Thumb, and Empirical Evidence," *Quantitative Finance*, 7, 5 (October 2007): 507–24.

are no longer correct. For example, delta of a call option is given by

$$\frac{\partial c_{BS}}{\partial S} + \frac{\partial c_{BS}}{\partial \sigma_{imp}} \frac{\partial \sigma_{imp}}{\partial S}$$

where c_{BS} is the Black-Scholes price of the option expressed as a function of the asset price S and the implied volatility σ_{imp} . Consider the impact of this formula on the delta of an equity call option. Volatility is a decreasing function of K/S . This means that the implied volatility increases as the asset price increases, so that

$$\frac{\partial \sigma_{imp}}{\partial S} > 0$$

As a result, delta is higher than that given by the Black–Scholes–Merton assumptions.

In practice, banks try to ensure that their exposure to the most commonly observed changes in the volatility surface is reasonably small. One technique for identifying these changes is principal components analysis, which we discuss in Chapter 21.

19.7 THE ROLE OF THE MODEL

How important is the option-pricing model if traders are prepared to use a different volatility for every option? It can be argued that the Black–Scholes–Merton model is no more than a sophisticated interpolation tool used by traders for ensuring that an option is priced consistently with the market prices of other actively traded options. If traders stopped using Black–Scholes–Merton and switched to another plausible model, then the volatility surface and the shape of the smile would change, but arguably the dollar prices quoted in the market would not change appreciably. Even delta, if calculated as outlined in the previous section, does not change too much as the model is changed.

Models have most effect on the pricing of derivatives when similar derivatives do not trade actively in the market. For example, the pricing of many of the nonstandard exotic derivatives we will discuss in later chapters is model-dependent.

Figure 19.5 Effect of a single large jump. The solid line is the true distribution; the dashed line is the lognormal distribution.

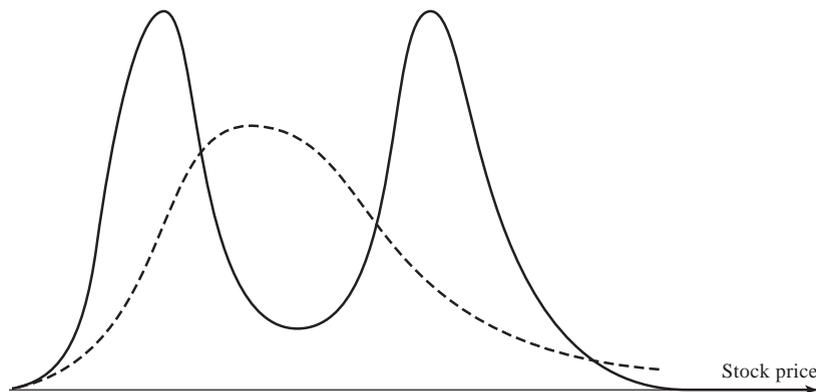
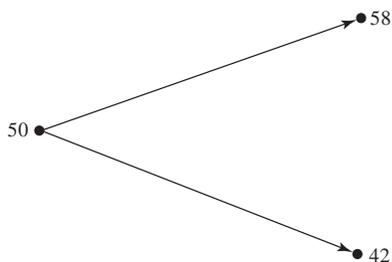


Figure 19.6 Change in stock price in 1 month.

19.8 WHEN A SINGLE LARGE JUMP IS ANTICIPATED

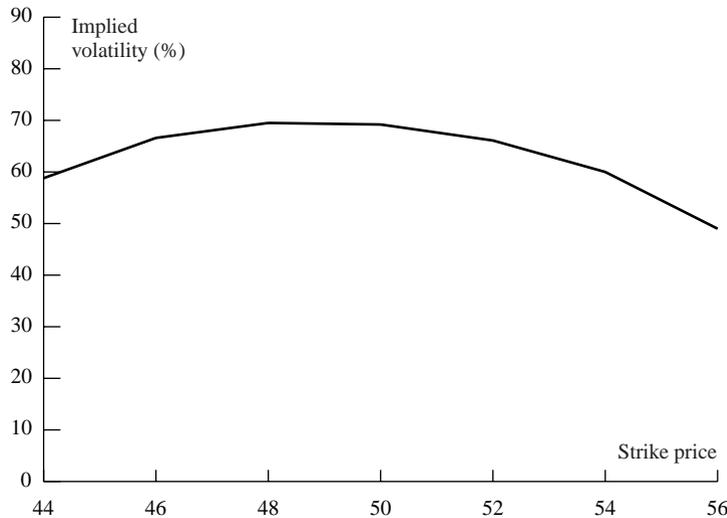
Let us now consider an example of how an unusual volatility smile might arise in equity markets. Suppose that a stock price is currently \$50 and an important news announcement due in a few days is expected either to increase the stock price by \$8 or to reduce it by \$8. (This announcement could concern the outcome of a takeover attempt or the verdict in an important lawsuit.) The probability distribution of the stock price in, say, 1 month might then consist of a mixture of two lognormal distributions, the first corresponding to favorable news, the second to unfavorable news. The situation is illustrated in Figure 19.5. The solid line shows the mixture-of-lognormals distribution for the stock price in 1 month; the dashed line shows a lognormal distribution with the same mean and standard deviation as this distribution.

The true probability distribution is bimodal (certainly not lognormal). One easy way to investigate the general effect of a bimodal stock price distribution is to consider the extreme case where the distribution is binomial. This is what we will now do.

Suppose that the stock price is currently \$50 and that it is known that in 1 month it will be either \$42 or \$58. Suppose further that the risk-free rate is 12% per annum. The situation is illustrated in Figure 19.6. Options can be valued using the binomial model from Chapter 12. In this case $u = 1.16$, $d = 0.84$, $a = 1.0101$, and $p = 0.5314$. The results from valuing a range of different options are shown in Table 19.3. The first

Table 19.3 Implied volatilities in situation where true distribution is binomial.

Strike price (\$)	Call price (\$)	Put price (\$)	Implied volatility (%)
42	8.42	0.00	0.0
44	7.37	0.93	58.8
46	6.31	1.86	66.6
48	5.26	2.78	69.5
50	4.21	3.71	69.2
52	3.16	4.64	66.1
54	2.10	5.57	60.0
56	1.05	6.50	49.0
58	0.00	7.42	0.0

Figure 19.7 Volatility smile for situation in Table 19.3.

column shows alternative strike prices; the second column shows prices of 1-month European call options; the third column shows the prices of one-month European put option prices; the fourth column shows implied volatilities. (As shown in Section 19.1, the implied volatility of a European put option is the same as that of a European call option when they have the same strike price and maturity.) Figure 19.7 displays the volatility smile from Table 19.3. It is actually a “frown” (the opposite of that observed for currencies) with volatilities declining as we move out of or into the money. The volatility implied from an option with a strike price of 50 will overprice an option with a strike price of 44 or 56.

SUMMARY

The Black–Scholes–Merton model and its extensions assume that the probability distribution of the underlying asset at any given future time is lognormal. This assumption is not the one made by traders. They assume the probability distribution of an equity price has a heavier left tail and a less heavy right tail than the lognormal distribution. They also assume that the probability distribution of an exchange rate has a heavier right tail and a heavier left tail than the lognormal distribution.

Traders use volatility smiles to allow for nonlognormality. The volatility smile defines the relationship between the implied volatility of an option and its strike price. For equity options, the volatility smile tends to be downward sloping. This means that out-of-the-money puts and in-the-money calls tend to have high implied volatilities whereas out-of-the-money calls and in-the-money puts tend to have low implied volatilities. For foreign currency options, the volatility smile is \cup -shaped. Both out-of-the-money and in-the-money options have higher implied volatilities than at-the-money options.

Often traders also use a volatility term structure. The implied volatility of an option then depends on its life. When volatility smiles and volatility term structures are

combined, they produce a volatility surface. This defines implied volatility as a function of both the strike price and the time to maturity.

FURTHER READING

- Bakshi, G., C. Cao, and Z. Chen. “Empirical Performance of Alternative Option Pricing Models,” *Journal of Finance*, 52, No. 5 (December 1997): 2004–49.
- Bates, D.S. “Post-’87 Crash Fears in the S&P Futures Market,” *Journal of Econometrics*, 94 (January/February 2000): 181–238.
- Daglish, T., J. Hull, and W. Suo. “Volatility Surfaces: Theory, Rules of Thumb, and Empirical Evidence,” *Quantitative Finance*, 7, 5 (2007), 507–24.
- Derman, E. “Regimes of Volatility,” *Risk*, April 1999: 55–59.
- Ederington, L. H., and W. Guan. “Why Are Those Options Smiling,” *Journal of Derivatives*, 10, 2 (2002): 9–34.
- Jackwerth, J.C., and M. Rubinstein. “Recovering Probability Distributions from Option Prices,” *Journal of Finance*, 51 (December 1996): 1611–31.
- Melick, W. R., and C.P. Thomas. “Recovering an Asset’s Implied Probability Density Function from Option Prices: An Application to Crude Oil during the Gulf Crisis,” *Journal of Financial and Quantitative Analysis*, 32, 1 (March 1997): 91–115.
- Reiswich, D., and U. Wystup. “FX Volatility Smile Construction,” Working Paper, Frankfurt School of Finance and Management, April 2010.
- Rubinstein, M. “Nonparametric Tests of Alternative Option Pricing Models Using All Reported Trades and Quotes on the 30 Most Active CBOE Option Classes from August 23, 1976, through August 31, 1978,” *Journal of Finance*, 40 (June 1985): 455–80.
- Xu, X., and S.J. Taylor. “The Term Structure of Volatility Implied by Foreign Exchange Options,” *Journal of Financial and Quantitative Analysis*, 29 (1994): 57–74.

Practice Questions (Answers in Solutions Manual)

- 19.1. What volatility smile is likely to be observed when:
 - (a) Both tails of the stock price distribution are less heavy than those of the lognormal distribution?
 - (b) The right tail is heavier, and the left tail is less heavy, than that of a lognormal distribution?
- 19.2. What volatility smile is observed for equities?
- 19.3. What volatility smile is likely to be caused by jumps in the underlying asset price? Is the pattern likely to be more pronounced for a 2-year option than for a 3-month option?
- 19.4. A European call and put option have the same strike price and time to maturity. The call has an implied volatility of 30% and the put has an implied volatility of 25%. What trades would you do?
- 19.5. Explain carefully why a distribution with a heavier left tail and less heavy right tail than the lognormal distribution gives rise to a downward sloping volatility smile.
- 19.6. The market price of a European call is \$3.00 and its price given by Black–Scholes–Merton model with a volatility of 30% is \$3.50. The price given by this Black–Scholes–Merton model for a European put option with the same strike price and time to maturity

- is \$1.00. What should the market price of the put option be? Explain the reasons for your answer.
- 19.7. Explain what is meant by “crashophobia.”
 - 19.8. A stock price is currently \$20. Tomorrow, news is expected to be announced that will either increase the price by \$5 or decrease the price by \$5. What are the problems in using Black–Scholes–Merton to value 1-month options on the stock?
 - 19.9. What volatility smile is likely to be observed for 6-month options when the volatility is uncertain and positively correlated to the stock price?
 - 19.10. What problems do you think would be encountered in testing a stock option pricing model empirically?
 - 19.11. Suppose that a central bank’s policy is to allow an exchange rate to fluctuate between 0.97 and 1.03. What pattern of implied volatilities for options on the exchange rate would you expect to see?
 - 19.12. Option traders sometimes refer to deep-out-of-the-money options as being options on volatility. Why do you think they do this?
 - 19.13. A European call option on a certain stock has a strike price of \$30, a time to maturity of 1 year, and an implied volatility of 30%. A European put option on the same stock has a strike price of \$30, a time to maturity of 1 year, and an implied volatility of 33%. What is the arbitrage opportunity open to a trader? Does the arbitrage work only when the lognormal assumption underlying Black–Scholes–Merton holds? Explain carefully the reasons for your answer.
 - 19.14. Suppose that the result of a major lawsuit affecting a company is due to be announced tomorrow. The company’s stock price is currently \$60. If the ruling is favorable to the company, the stock price is expected to jump to \$75. If it is unfavorable, the stock is expected to jump to \$50. What is the risk-neutral probability of a favorable ruling? Assume that the volatility of the company’s stock will be 25% for 6 months after the ruling if the ruling is favorable and 40% if it is unfavorable. Use DerivaGem to calculate the relationship between implied volatility and strike price for 6-month European options on the company today. The company does not pay dividends. Assume that the 6-month risk-free rate is 6%. Consider call options with strike prices of \$30, \$40, \$50, \$60, \$70, and \$80.
 - 19.15. An exchange rate is currently 0.8000. The volatility of the exchange rate is quoted as 12% and interest rates in the two countries are the same. Using the lognormal assumption, estimate the probability that the exchange rate in 3 months will be (a) less than 0.7000, (b) between 0.7000 and 0.7500, (c) between 0.7500 and 0.8000, (d) between 0.8000 and 0.8500, (e) between 0.8500 and 0.9000, and (f) greater than 0.9000. Based on the volatility smile usually observed in the market for exchange rates, which of these estimates would you expect to be too low and which would you expect to be too high?
 - 19.16. A stock price is \$40. A 6-month European call option on the stock with a strike price of \$30 has an implied volatility of 35%. A 6-month European call option on the stock with a strike price of \$50 has an implied volatility of 28%. The 6-month risk-free rate is 5% and no dividends are expected. Explain why the two implied volatilities are different. Use DerivaGem to calculate the prices of the two options. Use put–call parity to calculate the prices of 6-month European put options with strike prices of \$30 and \$50. Use DerivaGem to calculate the implied volatilities of these two put options.

- 19.17. “The Black–Scholes–Merton model is used by traders as an interpolation tool.” Discuss this view.
- 19.18. Using Table 19.2, calculate the implied volatility a trader would use for an 8-month option with $K/S_0 = 1.04$.

Further Questions

- 19.19. A company’s stock is selling for \$4. The company has no outstanding debt. Analysts consider the liquidation value of the company to be at least \$300,000 and there are 100,000 shares outstanding. What volatility smile would you expect to see?
- 19.20. A company is currently awaiting the outcome of a major lawsuit. This is expected to be known within 1 month. The stock price is currently \$20. If the outcome is positive, the stock price is expected to be \$24 at the end of 1 month. If the outcome is negative, it is expected to be \$18 at this time. The 1-month risk-free interest rate is 8% per annum.
- What is the risk-neutral probability of a positive outcome?
 - What are the values of 1-month call options with strike prices of \$19, \$20, \$21, \$22, and \$23?
 - Use DerivaGem to calculate a volatility smile for 1-month call options.
 - Verify that the same volatility smile is obtained for 1-month put options.
- 19.21. A futures price is currently \$40. The risk-free interest rate is 5%. Some news is expected tomorrow that will cause the volatility over the next 3 months to be either 10% or 30%. There is a 60% chance of the first outcome and a 40% chance of the second outcome. Use DerivaGem to calculate a volatility smile for 3-month options.
- 19.22. Data for a number of foreign currencies are provided on the author’s website:
<http://www.rotman.utoronto.ca/~hull/data>
Choose a currency and use the data to produce a table similar to Table 19.1.
- 19.23. Data for a number of stock indices are provided on the author’s website:
<http://www.rotman.utoronto.ca/~hull/data>
Choose an index and test whether a three-standard-deviation down movement happens more often than a three-standard-deviation up movement.
- 19.24. Consider a European call and a European put with the same strike price and time to maturity. Show that they change in value by the same amount when the volatility increases from a level σ_1 to a new level σ_2 within a short period of time. (*Hint*: Use put–call parity.)
- 19.25. An exchange rate is currently 1.0 and the implied volatilities of 6-month European options with strike prices 0.7, 0.8, 0.9, 1.0, 1.1, 1.2, 1.3 are 13%, 12%, 11%, 10%, 11%, 12%, 13%. The domestic and foreign risk-free rates are both 2.5%. Calculate the implied probability distribution using an approach similar to that used for Example 19A.1 in the appendix to this chapter. Compare it with the implied distribution where all the implied volatilities are 11.5%.
- 19.26. Using Table 19.2, calculate the implied volatility a trader would use for an 11-month option with $K/S_0 = 0.98$.

APPENDIX

DETERMINING IMPLIED RISK-NEUTRAL DISTRIBUTIONS FROM VOLATILITY SMILES

The price of a European call option on an asset with strike price K and maturity T is given by

$$c = e^{-rT} \int_{S_T=K}^{\infty} (S_T - K) g(S_T) dS_T$$

where r is the interest rate (assumed constant), S_T is the asset price at time T , and g is the risk-neutral probability density function of S_T . Differentiating once with respect to K gives

$$\frac{\partial c}{\partial K} = -e^{-rT} \int_{S_T=K}^{\infty} g(S_T) dS_T$$

Differentiating again with respect to K gives

$$\frac{\partial^2 c}{\partial K^2} = e^{-rT} g(K)$$

This shows that the probability density function g is given by

$$g(K) = e^{rT} \frac{\partial^2 c}{\partial K^2} \quad (19A.1)$$

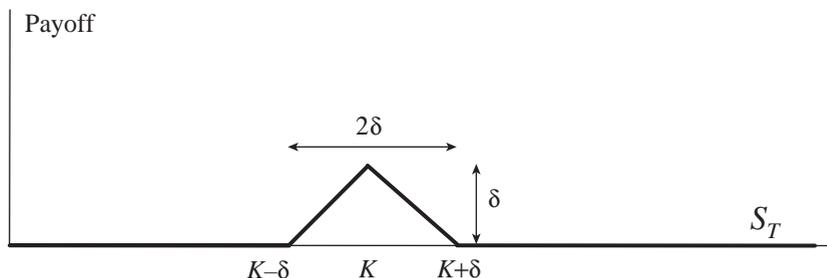
This result, which is from Breeden and Litzenberger (1978), allows risk-neutral probability distributions to be estimated from volatility smiles.⁹ Suppose that c_1 , c_2 , and c_3 are the prices of T -year European call options with strike prices of $K - \delta$, K , and $K + \delta$, respectively. Assuming δ is small, an estimate of $g(K)$, obtained by approximating the partial derivative in equation (19A.1), is

$$e^{rT} \frac{c_1 + c_3 - 2c_2}{\delta^2}$$

For another way of understanding this formula, suppose you set up a butterfly spread with strike prices $K - \delta$, K , and $K + \delta$, and maturity T . This means that you buy a call with strike price $K - \delta$, buy a call with strike price $K + \delta$, and sell two calls with strike price K . The value of your position is $c_1 + c_3 - 2c_2$. The value of the position can also be calculated by integrating the payoff over the risk-neutral probability distribution, $g(S_T)$, and discounting at the risk-free rate. The payoff is shown in Figure 19A.1. Since δ is small, we can assume that $g(S_T) = g(K)$ in the whole of the range $K - \delta < S_T < K + \delta$, where the payoff is nonzero. The area under the “spike” in Figure 19A.1 is $0.5 \times 2\delta \times \delta = \delta^2$. The value of the payoff (when δ is small) is therefore $e^{-rT} g(K) \delta^2$. It follows that

$$e^{-rT} g(K) \delta^2 = c_1 + c_3 - 2c_2$$

⁹ See D. T. Breeden and R. H. Litzenberger, “Prices of State-Contingent Claims Implicit in Option Prices,” *Journal of Business*, 51 (1978), 621–51.

Figure 19A.1 Payoff from butterfly spread.

which leads directly to

$$g(K) = e^{rT} \frac{c_1 + c_3 - 2c_2}{\delta^2} \quad (19A.2)$$

Example 19A.1

Suppose that the price of a non-dividend-paying stock is \$10, the risk-free interest rate is 3%, and the implied volatilities of 3-month European options with strike prices of \$6, \$7, \$8, \$9, \$10, \$11, \$12, \$13, \$14 are 30%, 29%, 28%, 27%, 26%, 25%, 24%, 23%, 22%, respectively. One way of applying the above results is as follows. Assume that $g(S_T)$ is constant between $S_T = 6$ and $S_T = 7$, constant between $S_T = 7$ and $S_T = 8$, and so on. Define:

$$\begin{aligned} g(S_T) &= g_1 & \text{for } 6 \leq S_T < 7 \\ g(S_T) &= g_2 & \text{for } 7 \leq S_T < 8 \\ g(S_T) &= g_3 & \text{for } 8 \leq S_T < 9 \\ g(S_T) &= g_4 & \text{for } 9 \leq S_T < 10 \\ g(S_T) &= g_5 & \text{for } 10 \leq S_T < 11 \\ g(S_T) &= g_6 & \text{for } 11 \leq S_T < 12 \\ g(S_T) &= g_7 & \text{for } 12 \leq S_T < 13 \\ g(S_T) &= g_8 & \text{for } 13 \leq S_T < 14 \end{aligned}$$

The value of g_1 can be calculated by interpolating to get the implied volatility for a 3-month option with a strike price of \$6.5 as 29.5%. This means that options with strike prices of \$6, \$6.5, and \$7 have implied volatilities of 30%, 29.5%, and 29%, respectively. From DerivaGem their prices are \$4.045, \$3.549, and \$3.055, respectively. Using equation (19A.2), with $K = 6.5$ and $\delta = 0.5$, gives

$$g_1 = \frac{e^{0.03 \times 0.25} (4.045 + 3.055 - 2 \times 3.549)}{0.5^2} = 0.0057$$

Similar calculations show that

$$\begin{aligned} g_2 &= 0.0444, & g_3 &= 0.1545, & g_4 &= 0.2781 \\ g_5 &= 0.2813, & g_6 &= 0.1659, & g_7 &= 0.0573, & g_8 &= 0.0113 \end{aligned}$$

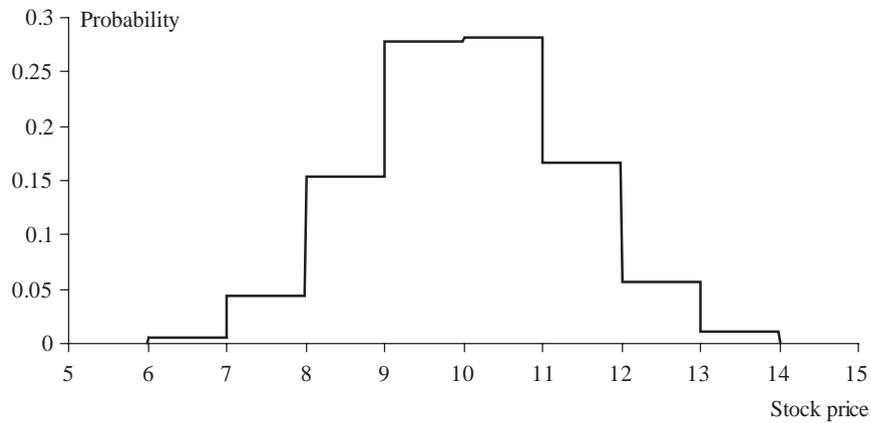
Figure 19.A.2 Implied probability distribution for Example 19A.1.

Figure 19A.2 displays the implied distribution. (Note that the area under the probability distribution is 0.9985. The probability that $S_T < 6$ or $S_T > 14$ is therefore 0.0015.) Although not obvious from Figure 19A.2, the implied distribution does have a heavier left tail and less heavy right tail than a lognormal distribution. For the lognormal distribution based on a single volatility of 26%, the probability of a stock price between \$6 and \$7 is 0.0031 (compared with 0.0057 in Figure 19A.2) and the probability of a stock price between \$13 and \$14 is 0.0167 (compared with 0.0113 in Figure 19A.2).

20

CHAPTER

Basic Numerical Procedures



This chapter discusses three numerical procedures for valuing derivatives when analytic results such as the Black–Scholes–Merton formulas do not exist. The first represents the asset price movements in the form of a tree and was introduced in Chapter 12. The second is Monte Carlo simulation, which we encountered briefly in Chapter 13 when stochastic processes were explained. The third involves finite difference methods.

Monte Carlo simulation is usually used for derivatives where the payoff is dependent on the history of the underlying variable or where there are several underlying variables. Trees and finite difference methods are usually used for American options and other derivatives where the holder has decisions to make prior to maturity. In addition to valuing a derivative, all the procedures can be used to calculate Greek letters such as delta, gamma, and vega.

The basic procedures discussed in this chapter can be used to handle most of the derivatives valuation problems encountered in practice. However, sometimes they have to be adapted to cope with particular situations, as will be explained in Chapter 26.

20.1 BINOMIAL TREES

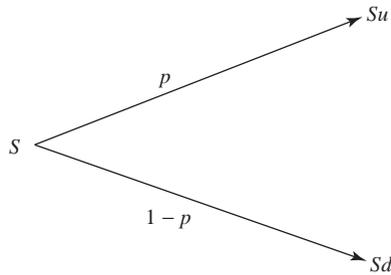
Binomial trees were introduced in Chapter 12. They can be used to value either European or American options. The Black–Scholes–Merton formulas and their extensions that were presented in Chapters 14, 16, and 17 provide analytic valuations for European options.¹ There are no analytic valuations for American options. Binomial trees are therefore most useful for valuing these types of options.²

As explained in Chapter 12, the binomial tree valuation approach involves dividing the life of the option into a large number of small time intervals of length Δt . It assumes that in each time interval the price of the underlying asset moves from its initial value of S to one of two new values, Su and Sd . The approach is illustrated in Figure 20.1. In

¹ The Black–Scholes–Merton formulas are based on the same set of assumptions as binomial trees. As shown in the appendix to Chapter 12, in the limit as the number of time steps is increased, the price given by a binomial tree for a European option converges to the Black–Scholes–Merton price.

² Some analytic approximations for valuing American options have been suggested. See, for example, Technical Note 8 at www.rotman.utoronto.ca/~hull/TechnicalNotes for a description of the quadratic approximation approach.

Figure 20.1 Asset price movements in time Δt under the binomial model.



general, $u > 1$ and $d < 1$. The movement from S to Su , therefore, is an “up” movement and the movement from S to Sd is a “down” movement. The probability of an up movement will be denoted by p . The probability of a down movement is $1 - p$.

Risk-Neutral Valuation

The risk-neutral valuation principle, explained in Chapters 12 and 14, states that an option (or other derivative) can be valued on the assumption that the world is risk neutral. This means that for valuation purposes we can use the following procedure:

1. Assume that the expected return from all traded assets is the risk-free interest rate.
2. Value payoffs from the derivative by calculating their expected values and discounting at the risk-free interest rate.

This principle underlies the way trees are used.

Determination of p , u , and d

The parameters p , u , and d must give correct values for the mean and variance of asset price changes during a time interval of length Δt . Because we are working in a risk-neutral world, the expected return from the asset is the risk-free interest rate, r . Suppose that the asset provides a yield of q . The expected return in the form of capital gains must be $r - q$. This means that the expected value of the asset price at the end of a time interval of length Δt must be $Se^{(r-q)\Delta t}$, where S is the asset price at the beginning of the time interval. To match the mean return with the tree, we therefore need

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

or

$$e^{(r-q)\Delta t} = pu + (1 - p)d \quad (20.1)$$

The variance of a variable Q is defined as $E(Q^2) - [E(Q)]^2$. Defining R as the percentage change in the asset price in time Δt , there is a probability p that $1 + R$ is u and a probability $1 - p$ that it is d . Using equation (20.1), it follows that the variance of $1 + R$ is

$$pu^2 + (1 - p)d^2 - e^{2(r-q)\Delta t}$$

Since adding a constant to a variable makes no difference to its variance, the variance of $1 + R$ is the same as the variance of R . As explained in Section 14.4, this is $\sigma^2 \Delta t$.

Hence,

$$pu^2 + (1 - p)d^2 - e^{2(r-q)\Delta t} = \sigma^2 \Delta t$$

From equation (20.1), $e^{(r-q)\Delta t}(u + d) = pu^2 + (1 - p)d^2 + ud$, so that

$$e^{(r-q)\Delta t}(u + d) - ud - e^{2(r-q)\Delta t} = \sigma^2 \Delta t \quad (20.2)$$

Equations (20.1) and (20.2) impose two conditions on p , u , and d . A third condition used by Cox, Ross, and Rubinstein (1979) is³

$$u = 1/d \quad (20.3)$$

A solution to equations (20.1) to (20.3), when terms of higher order than Δt are ignored, is⁴

$$p = \frac{a - d}{u - d} \quad (20.4)$$

$$u = e^{\sigma\sqrt{\Delta t}} \quad (20.5)$$

$$d = e^{-\sigma\sqrt{\Delta t}} \quad (20.6)$$

where

$$a = e^{(r-q)\Delta t} \quad (20.7)$$

The variable a is sometimes referred to as the *growth factor*. Equations (20.4) to (20.7) are consistent with the formulas in Sections 12.9 and 12.11.

Tree of Asset Prices

Figure 20.2 shows the complete tree of asset prices that is considered when the binomial model is used with five time steps. At time zero, the asset price, S_0 , is known. At time Δt , there are two possible asset prices, S_0u and S_0d ; at time $2\Delta t$, there are three possible asset prices, S_0u^2 , S_0 , and S_0d^2 ; and so on. In general, at time $i\Delta t$, we consider $i + 1$ asset prices. These are

$$S_0u^j d^{i-j}, \quad j = 0, 1, \dots, i$$

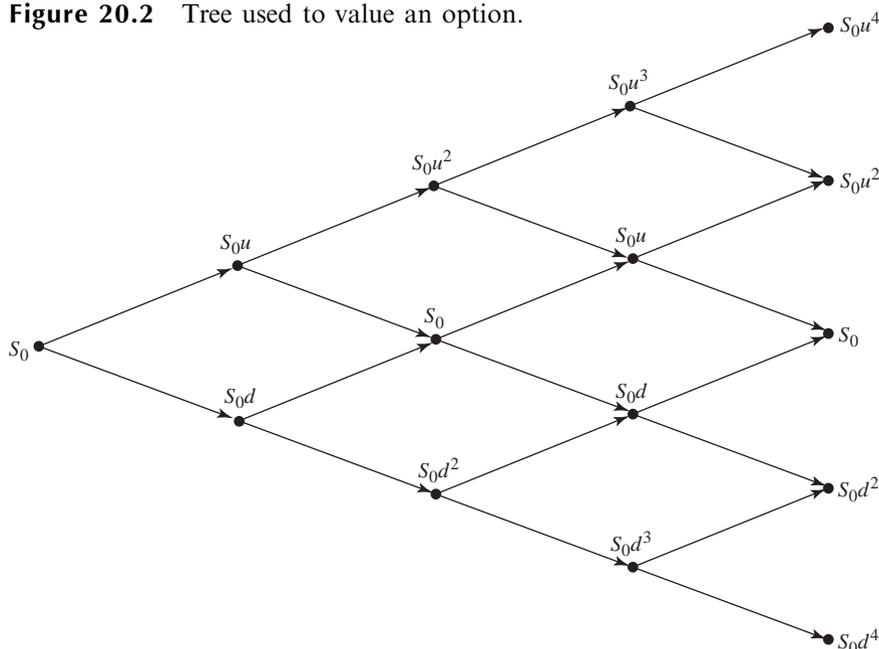
Note that the relationship $u = 1/d$ is used in computing the asset price at each node of the tree in Figure 20.2. For example, the asset price when $j = 2$ and $i = 3$ is $S_0u^2d = S_0u$. Note also that the tree recombines in the sense that an up movement followed by a down movement leads to the same asset price as a down movement followed by an up movement.

Working Backward through the Tree

Options are evaluated by starting at the end of the tree (time T) and working backward. The value of the option is known at time T . For example, a put option is worth $\max(K - S_T, 0)$ and a call option is worth $\max(S_T - K, 0)$, where S_T is the asset price at

³ See J.C. Cox, S.A. Ross, and M. Rubinstein, "Option Pricing: A Simplified Approach," *Journal of Financial Economics*, 7 (October 1979), 229–63.

⁴ To see this, we note that equations (20.4) and (20.7) satisfy the conditions in equations (20.1) and (20.3) exactly. The exponential function e^x can be expanded as $1 + x + x^2/2 + \dots$. When terms of higher order than Δt are ignored, equation (20.5) implies that $u = 1 + \sigma\sqrt{\Delta t} + \frac{1}{2}\sigma^2\Delta t$ and equation (20.6) implies that $d = 1 - \sigma\sqrt{\Delta t} + \frac{1}{2}\sigma^2\Delta t$. Also, $e^{(r-q)\Delta t} = 1 + (r - q)\Delta t$ and $e^{2(r-q)\Delta t} = 1 + 2(r - q)\Delta t$. By substitution, we see that equation (20.2) is satisfied when terms of higher order than Δt are ignored.

Figure 20.2 Tree used to value an option.

time T and K is the strike price. Because a risk-neutral world is being assumed, the value at each node at time $T - \Delta t$ can be calculated as the expected value at time T discounted at rate r for a time period Δt . Similarly, the value at each node at time $T - 2\Delta t$ can be calculated as the expected value at time $T - \Delta t$ discounted for a time period Δt at rate r , and so on. If the option is American, it is necessary to check at each node to see whether early exercise is preferable to holding the option for a further time period Δt . Eventually, by working back through all the nodes, we are able to obtain the value of the option at time zero.

Example 20.1

Consider a 5-month American put option on a non-dividend-paying stock when the stock price is \$50, the strike price is \$50, the risk-free interest rate is 10% per annum, and the volatility is 40% per annum. With our usual notation, this means that $S_0 = 50$, $K = 50$, $r = 0.10$, $\sigma = 0.40$, $T = 0.4167$, and $q = 0$. Suppose that we divide the life of the option into five intervals of length 1 month ($= 0.0833$ year) for the purposes of constructing a binomial tree. Then $\Delta t = 0.0833$ and using equations (20.4) to (20.7) gives

$$u = e^{\sigma\sqrt{\Delta t}} = 1.1224, \quad d = e^{-\sigma\sqrt{\Delta t}} = 0.8909, \quad a = e^{r\Delta t} = 1.0084$$

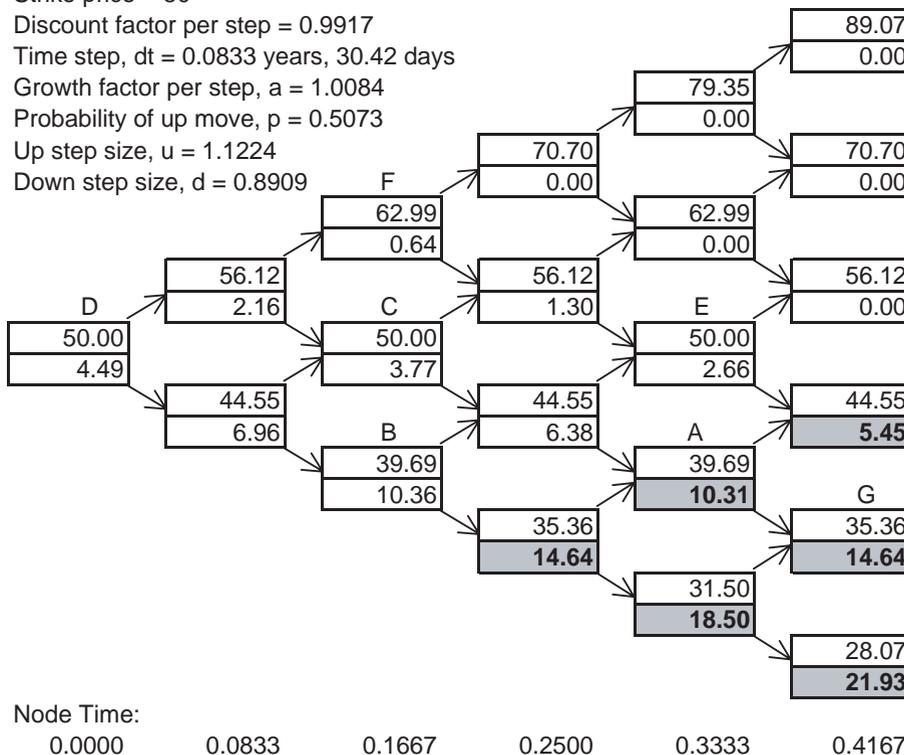
$$p = \frac{a - d}{u - d} = 0.5073, \quad 1 - p = 0.4927$$

Figure 20.3 shows the binomial tree produced by DerivaGem. At each node there are two numbers. The top one shows the stock price at the node; the lower one shows the value of the option at the node. The probability of an up movement is always 0.5073; the probability of a down movement is always 0.4927.

Figure 20.3 Binomial tree from DerivaGem for American put on non-dividend-paying stock (Example 20.1).

At each node:
 Upper value = Underlying Asset Price
 Lower value = Option Price
 Shading indicates where option is exercised

Strike price = 50
 Discount factor per step = 0.9917
 Time step, $dt = 0.0833$ years, 30.42 days
 Growth factor per step, $a = 1.0084$
 Probability of up move, $p = 0.5073$
 Up step size, $u = 1.1224$
 Down step size, $d = 0.8909$



The stock price at the j th node ($j = 0, 1, \dots, i$) at time $i \Delta t$ ($i = 0, 1, \dots, 5$) is calculated as $S_0 u^j d^{i-j}$. For example, the stock price at node A ($i = 4, j = 1$) (i.e., the second node up at the end of the fourth time step) is $50 \times 1.1224 \times 0.8909^3 = \39.69 . The option prices at the final nodes are calculated as $\max(K - S_T, 0)$. For example, the option price at node G is $50.00 - 35.36 = 14.64$. The option prices at the penultimate nodes are calculated from the option prices at the final nodes. First, we assume no exercise of the option at the nodes. This means that the option price is calculated as the present value of the expected option price one time step later. For example, at node E, the option price is calculated as

$$(0.5073 \times 0 + 0.4927 \times 5.45)e^{-0.10 \times 0.0833} = 2.66$$

whereas at node A it is calculated as

$$(0.5073 \times 5.45 + 0.4927 \times 14.64)e^{-0.10 \times 0.0833} = 9.90$$

We then check to see if early exercise is preferable to waiting. At node E, early exercise would give a value for the option of zero because both the stock price and strike price are \$50. Clearly it is best to wait. The correct value for the option at node E is therefore \$2.66. At node A, it is a different story. If the option is exercised, it is worth \$50.00 – \$39.69, or \$10.31. This is more than \$9.90. If node A is reached, then the option should be exercised and the correct value for the option at node A is \$10.31.

Option prices at earlier nodes are calculated in a similar way. Note that it is not always best to exercise an option early when it is in the money. Consider node B. If the option is exercised, it is worth \$50.00 – \$39.69, or \$10.31. However, if it is not exercised, it is worth

$$(0.5073 \times 6.38 + 0.4927 \times 14.64)e^{-0.10 \times 0.0833} = 10.36$$

The option should, therefore, not be exercised at this node, and the correct option value at the node is \$10.36.

Working back through the tree, the value of the option at the initial node is \$4.49. This is our numerical estimate for the option's current value. In practice, a smaller value of Δt , and many more nodes, would be used. DerivaGem shows that with 30, 50, 100, and 500 time steps we get values for the option of 4.263, 4.272, 4.278, and 4.283.

Expressing the Approach Algebraically

Suppose that the life of an American option is divided into N subintervals of length Δt . We will refer to the j th node at time $i \Delta t$ as the (i, j) node, where $0 \leq i \leq N$ and $0 \leq j \leq i$. Define $f_{i,j}$ as the value of the option at the (i, j) node. The price of the underlying asset at the (i, j) node is $S_0 u^j d^{i-j}$. If the option is a call, its value at time T (the expiration date) is $\max(S_T - K, 0)$, so that

$$f_{N,j} = \max(S_0 u^j d^{N-j} - K, 0), \quad j = 0, 1, \dots, N$$

If the option is a put, its value at time T is $\max(K - S_T, 0)$, so that

$$f_{N,j} = \max(K - S_0 u^j d^{N-j}, 0), \quad j = 0, 1, \dots, N$$

There is a probability p of moving from the (i, j) node at time $i \Delta t$ to the $(i + 1, j + 1)$ node at time $(i + 1) \Delta t$, and a probability $1 - p$ of moving from the (i, j) node at time $i \Delta t$ to the $(i + 1, j)$ node at time $(i + 1) \Delta t$. Assuming no early exercise, risk-neutral valuation gives

$$f_{i,j} = e^{-r\Delta t} [p f_{i+1,j+1} + (1-p) f_{i+1,j}]$$

for $0 \leq i \leq N - 1$ and $0 \leq j \leq i$. To take account of early exercise, this value for $f_{i,j}$ must be compared with the option's intrinsic value, so that for a call

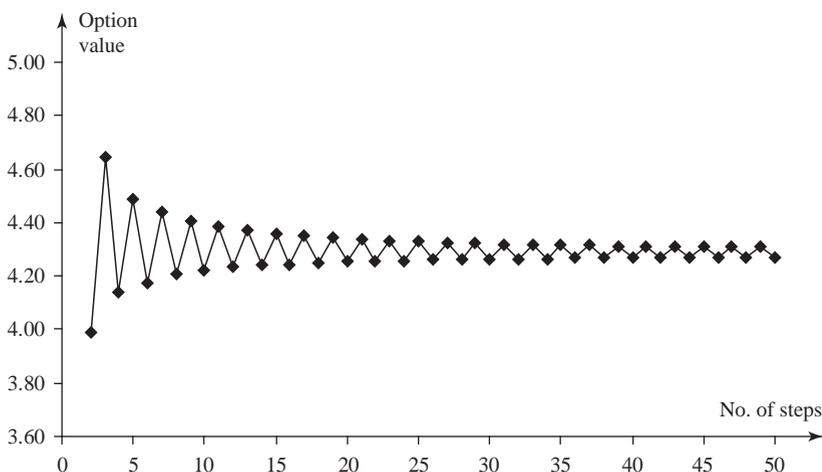
$$f_{i,j} = \max\{S_0 u^j d^{i-j} - K, e^{-r\Delta t} [p f_{i+1,j+1} + (1-p) f_{i+1,j}]\}$$

and for a put

$$f_{i,j} = \max\{K - S_0 u^j d^{i-j}, e^{-r\Delta t} [p f_{i+1,j+1} + (1-p) f_{i+1,j}]\}$$

Note that, because the calculations start at time T and work backward, the value at

Figure 20.4 Convergence of the price of the option in Example 20.1 calculated from the DerivaGem Application Builder functions.



time $i \Delta t$ captures not only the effect of early exercise possibilities at time $i \Delta t$, but also the effect of early exercise at subsequent times.

In the limit as Δt tends to zero, an exact value for the American put is obtained. In practice, $N = 30$ usually gives reasonable results. Figure 20.4 shows the convergence of the option price in Example 20.1. This figure was calculated using the Application Builder functions provided with the DerivaGem software (see Sample Application A).

Estimating Delta and Other Greek Letters

It will be recalled that the delta (Δ) of an option is the rate of change of its price with respect to the underlying stock price. It can be calculated as

$$\frac{\Delta f}{\Delta S}$$

where ΔS is a small change in the asset price and Δf is the corresponding small change in the option price. At time Δt , we have an estimate $f_{1,1}$ for the option price when the asset price is $S_0 u$ and an estimate $f_{1,0}$ for the option price when the asset price is $S_0 d$. This means that, when $\Delta S = S_0 u - S_0 d$, $\Delta f = f_{1,1} - f_{1,0}$. Therefore an estimate of delta at time Δt is

$$\Delta = \frac{f_{1,1} - f_{1,0}}{S_0 u - S_0 d} \quad (20.8)$$

To determine gamma (Γ), note that we have two estimates of Δ at time $2\Delta t$. When $S = (S_0 u^2 + S_0)/2$ (halfway between the second and third node), delta is $(f_{2,2} - f_{2,1})/(S_0 u^2 - S_0)$; when $S = (S_0 + S_0 d^2)/2$ (halfway between the first and second node), delta is $(f_{2,1} - f_{2,0})/(S_0 - S_0 d^2)$. The difference between the two values of S is h , where

$$h = 0.5(S_0 u^2 - S_0 d^2)$$

Gamma is the change in delta divided by h :

$$\Gamma = \frac{[(f_{2,2} - f_{2,1})/(S_0u^2 - S_0)] - [(f_{2,1} - f_{2,0})/(S_0 - S_0d^2)]}{h} \quad (20.9)$$

These procedures provide estimates of delta at time Δt and of gamma at time $2\Delta t$. In practice, they are usually used as estimates of delta and gamma at time zero as well.⁵

A further hedge parameter that can be obtained directly from the tree is theta (Θ). This is the rate of change of the option price with time when all else is kept constant. For an asset price of S_0 , the value of the option at time zero is $f_{0,0}$ and at time $2\Delta t$ it is $f_{2,1}$. An estimate of theta is therefore

$$\Theta = \frac{f_{2,1} - f_{0,0}}{2\Delta t} \quad (20.10)$$

Vega can be calculated by making a small change, $\Delta\sigma$, in the volatility and constructing a new tree to obtain a new value of the option. (The number of time steps should be kept the same.) The estimate of vega is

$$\mathcal{V} = \frac{f^* - f}{\Delta\sigma}$$

where f and f^* are the estimates of the option price from the original and the new tree, respectively. Rho can be calculated similarly.

Example 20.2

Consider again Example 20.1. From Figure 20.3, $f_{1,0} = 6.96$ and $f_{1,1} = 2.16$. Equation (20.8) gives an estimate for delta of

$$\frac{2.16 - 6.96}{56.12 - 44.55} = -0.41$$

From equation (20.9), an estimate of the gamma of the option can be obtained from the values at nodes B, C, and F as

$$\frac{[(0.64 - 3.77)/(62.99 - 50.00)] - [(3.77 - 10.36)/(50.00 - 39.69)]}{11.65} = 0.03$$

From equation (20.10), an estimate of the theta of the option can be obtained from the values at nodes D and C as

$$\frac{3.77 - 4.49}{0.1667} = -4.3 \quad \text{per year}$$

or -0.012 per calendar day. These are only rough estimates. They become progressively better as the number of time steps on the tree is increased. Using 50 time steps, DerivaGem provides estimates of -0.415 , 0.034 , and -0.0117 for delta, gamma, and theta, respectively. By making small changes to parameters and recomputing values, vega and rho are estimated as 0.123 and -0.072 , respectively.

⁵ If slightly more accuracy is required for delta and gamma, we can start the binomial tree at time $-2\Delta t$ and assume that the stock price is S_0 at this time. This leads to the option price being calculated for three different stock prices at time zero.

20.2 USING THE BINOMIAL TREE FOR OPTIONS ON INDICES, CURRENCIES, AND FUTURES CONTRACTS

As explained in Chapters 12, 16 and 17, stock indices, currencies, and futures contracts can, for the purposes of option valuation, be considered as assets providing known yields. For a stock index, the relevant yield is the dividend yield on the stock portfolio underlying the index; in the case of a currency, it is the foreign risk-free interest rate; in the case of a futures contract, it is the domestic risk-free interest rate. The binomial tree approach can therefore be used to value options on stock indices, currencies, and futures contracts provided that q in equation (20.7) is interpreted appropriately.

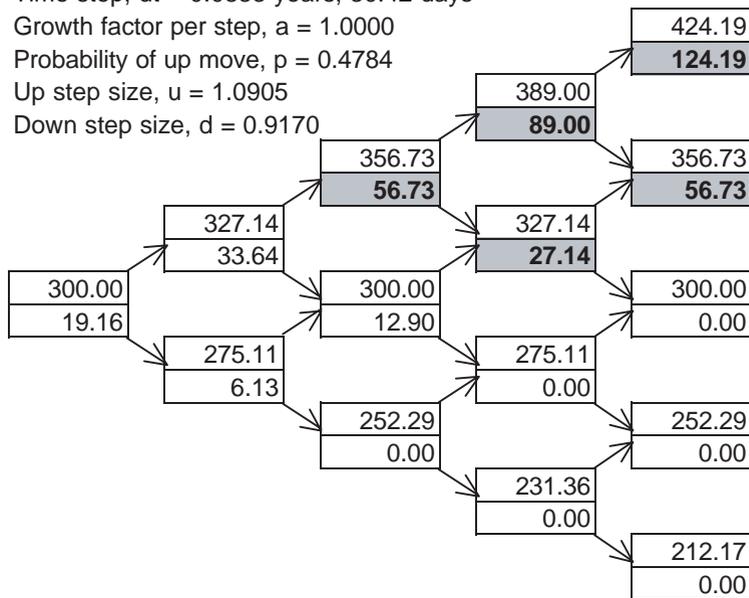
Example 20.3

Consider a 4-month American call option on index futures where the current futures price is 300, the exercise price is 300, the risk-free interest rate is 8% per

Figure 20.5 Binomial tree produced by DerivaGem for American call option on an index futures contract (Example 20.3).

At each node:
 Upper value = Underlying Asset Price
 Lower value = Option Price
 Shading indicates where option is exercised

Strike price = 300
 Discount factor per step = 0.9934
 Time step, $dt = 0.0833$ years, 30.42 days
 Growth factor per step, $a = 1.0000$
 Probability of up move, $p = 0.4784$
 Up step size, $u = 1.0905$
 Down step size, $d = 0.9170$



Node Time:
 0.0000 0.0833 0.1667 0.2500 0.3333

annum, and the volatility of the index is 30% per annum. The life of the option is divided into four 1-month periods for the purposes of constructing the tree. In this case, $F_0 = 300$, $K = 300$, $r = 0.08$, $\sigma = 0.3$, $T = 0.3333$, and $\Delta t = 0.0833$. Because a futures contract is analogous to a stock paying dividends at a rate r , q should be set equal to r in equation (20.7). This gives $a = 1$. The other parameters necessary to construct the tree are

$$u = e^{\sigma\sqrt{\Delta t}} = 1.0905, \quad d = 1/u = 0.9170$$

$$p = \frac{a - d}{u - d} = 0.4784, \quad 1 - p = 0.5216$$

The tree, as produced by DerivaGem, is shown in Figure 20.5. (The upper number is the futures price; the lower number is the option price.) The estimated value of the option is 20.16. More accuracy is obtained using more steps. With 50 time steps, DerivaGem gives a value of 20.18; with 100 time steps it gives 20.22.

Figure 20.6 Binomial tree produced by DerivaGem for American put option on a currency (Example 20.4).

At each node:

Upper value = Underlying Asset Price

Lower value = Option Price

Shading indicates where option is exercised

Strike price = 1.6

Discount factor per step = 0.9802

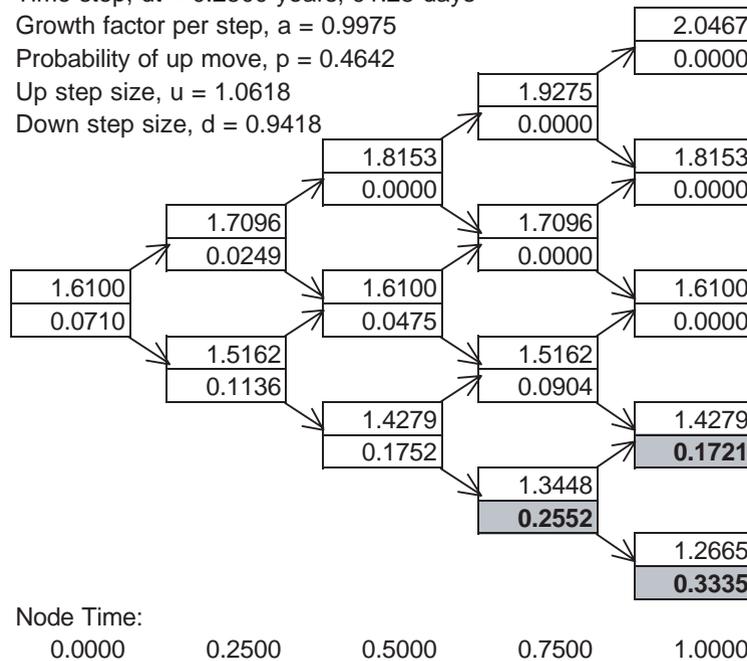
Time step, $dt = 0.2500$ years, 91.25 days

Growth factor per step, $a = 0.9975$

Probability of up move, $p = 0.4642$

Up step size, $u = 1.0618$

Down step size, $d = 0.9418$



Example 20.4

Consider a 1-year American put option on the British pound (GBP). The current exchange rate (USD per GBP) is 1.6100, the strike price is 1.6000, the US risk-free interest rate is 8% per annum, the sterling risk-free interest rate is 9% per annum, and the volatility of the sterling exchange rate is 12% per annum. In this case, $S_0 = 1.61$, $K = 1.60$, $r = 0.08$, $r_f = 0.09$, $\sigma = 0.12$, and $T = 1.0$. The life of the option is divided into four 3-month periods for the purposes of constructing the tree, so that $\Delta t = 0.25$. In this case, $q = r_f$ and equation (20.7) gives

$$a = e^{(0.08-0.09) \times 0.25} = 0.9975$$

The other parameters necessary to construct the tree are

$$u = e^{\sigma\sqrt{\Delta t}} = 1.0618, \quad d = 1/u = 0.9418 \quad p = \frac{a-d}{u-d} = 0.4642, \quad 1-p = 0.5358$$

The tree, as produced by DerivaGem, is shown in Figure 20.6. (The upper number is the exchange rate; the lower number is the option price.) The estimated value of the option is \$0.0710. (Using 50 time steps, DerivaGem gives the value of the option as 0.0738; with 100 time steps it also gives 0.0738.)

20.3 BINOMIAL MODEL FOR A DIVIDEND-PAYING STOCK

We now move on to the more tricky issue of how the binomial model can be used for a dividend-paying stock. As in Chapter 14, the word “dividend” will, for the purposes of our discussion, be used to refer to the reduction in the stock price on the ex-dividend date as a result of the dividend.

Known Dividend Yield

For long-life stock options, it is sometimes assumed for convenience that there is a known continuous dividend yield of q on the stock. The options can then be valued in the same way as options on a stock index.

For more accuracy, known dividend yields can be assumed to be paid discretely. Suppose that there is a single dividend, and the dividend yield (i.e., the dividend as a percentage of the stock price) is known. The parameters u , d , and p can be calculated as though no dividends are expected. If the time $i \Delta t$ is prior to the stock going ex-dividend, the nodes on the tree correspond to stock prices

$$S_0 u^j d^{i-j}, \quad j = 0, 1, \dots, i$$

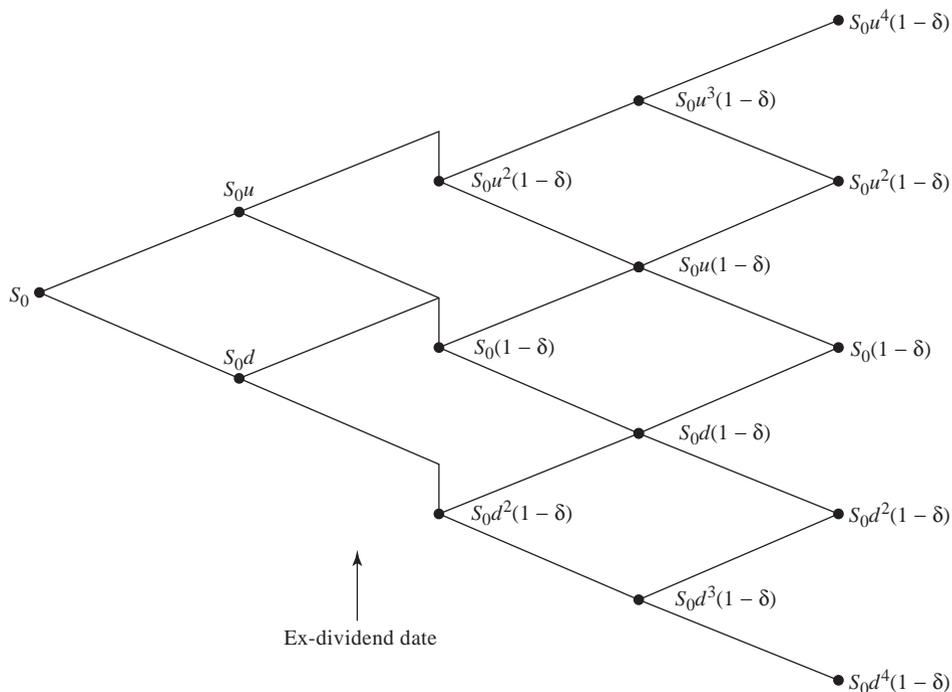
If the time $i \Delta t$ is after the stock goes ex-dividend, the nodes correspond to stock prices

$$S_0(1-\delta)u^j d^{i-j}, \quad j = 0, 1, \dots, i$$

where δ is the dividend yield. Several known dividend yields during the life of an option can be dealt with similarly. If δ_i is the total dividend yield associated with all ex-dividend dates between time zero and time $i \Delta t$, the nodes at time $i \Delta t$ correspond to stock prices

$$S_0(1-\delta_i)u^j d^{i-j}$$

Figure 20.7 Tree when stock pays a known dividend yield at one particular time.



Known Dollar Dividend

In some situations, particularly when the life of the option is short, the most realistic assumption is that the dollar amount of the dividend rather than the dividend yield is known in advance. If the volatility of the stock, σ , is assumed constant, the tree then takes the form shown in Figure 20.8. It does not recombine, which means that the number of nodes that have to be evaluated is liable to become very large. Suppose that there is only one dividend, that the ex-dividend date, τ , is between $k \Delta t$ and $(k + 1) \Delta t$, and that the dollar amount of the dividend is D . When $i \leq k$, the nodes on the tree at time $i \Delta t$ correspond to stock prices

$$S_0 u^j d^{i-j}, \quad j = 0, 1, 2, \dots, i$$

as before. When $i = k + 1$, the nodes on the tree correspond to stock prices

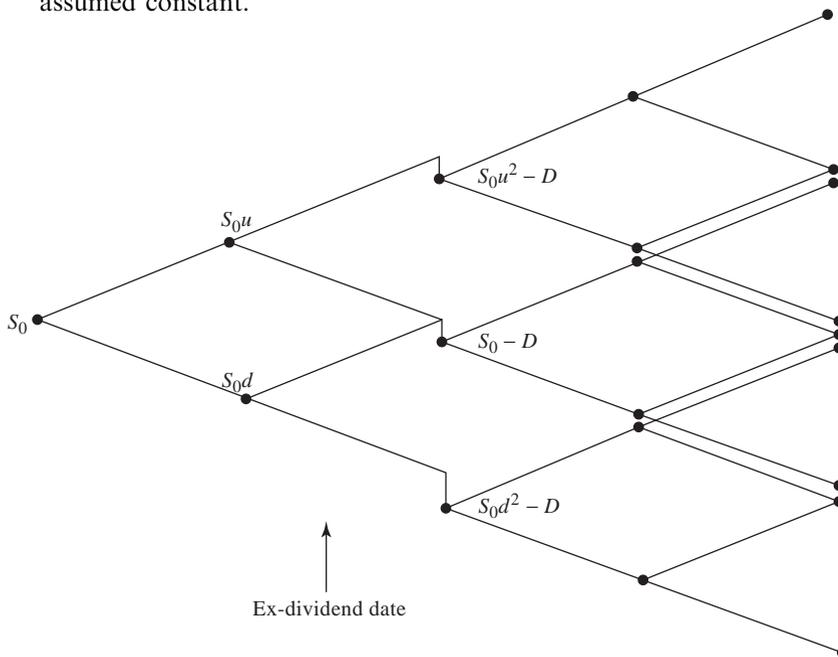
$$S_0 u^j d^{i-j} - D, \quad j = 0, 1, 2, \dots, i$$

When $i = k + 2$, the nodes on the tree correspond to stock prices

$$(S_0 u^j d^{i-1-j} - D)u \quad \text{and} \quad (S_0 u^j d^{i-1-j} - D)d$$

for $j = 0, 1, 2, \dots, i - 1$, so that there are $2i$ rather than $i + 1$ nodes. When $i = k + m$, there are $m(k + 2)$ rather than $k + m + 1$ nodes. The number of nodes expands even faster when there are several ex-dividend dates during the option's life.

Figure 20.8 Tree when dollar amount of dividend is assumed known and volatility is assumed constant.



The node-proliferation problem can be solved by assuming, as in the analysis of European options in Section 14.12, that the stock price has two components: a part that is uncertain and a part that is the present value of all future dividends during the life of the option. Suppose that there is only one ex-dividend date, τ , during the life of the option and that $k \Delta t \leq \tau \leq (k + 1) \Delta t$. The value of the uncertain component, S^* , at time $i \Delta t$ is given by

$$S^* = S \quad \text{when } i \Delta t > \tau$$

and

$$S^* = S - De^{-r(\tau - i \Delta t)} \quad \text{when } i \Delta t \leq \tau$$

where D is the dividend. Define σ^* as the volatility of S^* and assume that σ^* is constant.⁶ The parameters p , u , and d can be calculated from equations (20.4), (20.5), (20.6), and (20.7) with σ replaced by σ^* and a tree can be constructed in the usual way to model S^* . By adding to the stock price at each node, the present value of future dividends (if any), the tree can be converted into another tree that models S . Suppose that S_0^* is the value of S^* at time zero. At time $i \Delta t$, the nodes on this tree correspond to the stock prices

$$S_0^* u^j d^{i-j} + De^{-r(\tau - i \Delta t)}, \quad j = 0, 1, \dots, i$$

⁶ As mentioned in footnote 12 of Chapter 14, σ^* is greater than σ , the volatility of S . In practice, the use of a term structure of implied volatilities avoids the need for analysts to distinguish between σ and σ^* .

when $i \Delta t < \tau$ and

$$S_0^* u^j d^{i-j}, \quad j = 0, 1, \dots, i$$

when $i \Delta t > \tau$. This approach, which has the advantage of being consistent with the approach for European options in Section 14.12, succeeds in achieving a situation where the tree recombines so that there are $i + 1$ nodes at time $i \Delta t$. It can be generalized in a straightforward way to deal with the situation where there are several dividends.⁷

Example 20.5

Consider a 5-month American put option on a stock that is expected to pay a single dividend of \$2.06 during the life of the option. The initial stock price is \$52, the strike price is \$50, the risk-free interest rate is 10% per annum, the volatility is 40% per annum, and the ex-dividend date is in $3\frac{1}{2}$ months.

We first construct a tree to model S^* , the stock price less the present value of future dividends during the life of the option. At time zero, the present value of the dividend is

$$2.06 \times e^{-0.2917 \times 0.1} = 2.00$$

The initial value of S^* is therefore 50.00. If we assume that the 40% per annum volatility refers to S^* , then Figure 20.3 provides a binomial tree for S^* . (This is because S^* has the same initial value and volatility as the stock price that Figure 20.3 was based upon.) Adding the present value of the dividend at each node leads to Figure 20.9, which is a binomial model for S . The probabilities at each node are, as in Figure 20.3, 0.5073 for an up movement and 0.4927 for a down movement. Working back through the tree in the usual way gives the option price as \$4.44. (Using 50 time steps, DerivaGem gives a value for the option of 4.202; using 100 steps it gives 4.212.)

Control Variate Technique

A technique known as the *control variate technique* can improve the accuracy of the pricing of an American option.⁸ This involves using the same tree to calculate the value of both the American option, f_A , and the corresponding European option, f_E . The Black–Scholes–Merton price of the European option, f_{BS} , is also calculated. The error when the tree is used to price the European option, $f_{BS} - f_E$, is assumed equal to the error when the tree is used to price the American option. This gives the estimate of the price of the American option as

$$f_A + (f_{BS} - f_E)$$

To illustrate this approach, Figure 20.10 values the option in Figure 20.3 on the assumption that it is European. The price obtained, f_E , is \$4.32. From the Black–Scholes–Merton formula, the true European price of the option, f_{BS} , is \$4.08. The

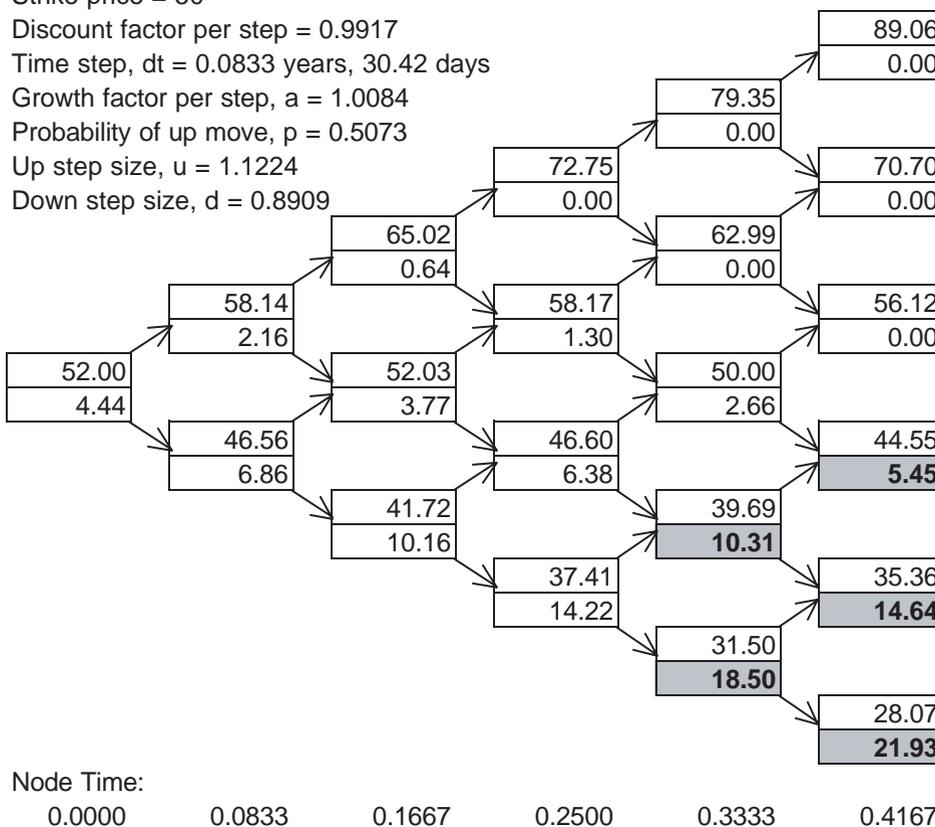
⁷ For long-life options, where there are many dividends, the dividends are less easy to predict and the present value of the dividends becomes a significant part of S_0 . It is then often more appropriate to assume a known dividend yield.

⁸ See J. Hull and A. White, “The Use of the Control Variate Technique in Option Pricing,” *Journal of Financial and Quantitative Analysis*, 23 (September 1988): 237–51.

Figure 20.9 Tree produced by DerivaGem for Example 20.5.

At each node:
 Upper value = Underlying Asset Price
 Lower value = Option Price
 Shading indicates where option is exercised

Strike price = 50
 Discount factor per step = 0.9917
 Time step, dt = 0.0833 years, 30.42 days
 Growth factor per step, a = 1.0084
 Probability of up move, p = 0.5073
 Up step size, u = 1.1224
 Down step size, d = 0.8909



estimate of the American price in Figure 20.3, f_A , is \$4.49. The control variate estimate of the American price, therefore, is

$$4.49 + (4.08 - 4.32) = 4.25$$

A good estimate of the American price, calculated using 100 steps, is 4.278. The control variate approach does, therefore, produce a considerable improvement over the basic tree estimate of 4.49 in this case.

The control variate technique in effect involves using the tree to calculate the difference between the European and the American price rather than the American price itself. We give a further application of the control variate technique when we discuss Monte Carlo simulation later in the chapter.

Figure 20.10 Tree, as produced by DerivaGem, for European version of option in Figure 20.3. At each node, the upper number is the stock price, and the lower number is the option price.

At each node:

Upper value = Underlying Asset Price

Lower value = Option Price

Shading indicates where option is exercised

Strike price = 50

Discount factor per step = 0.9917

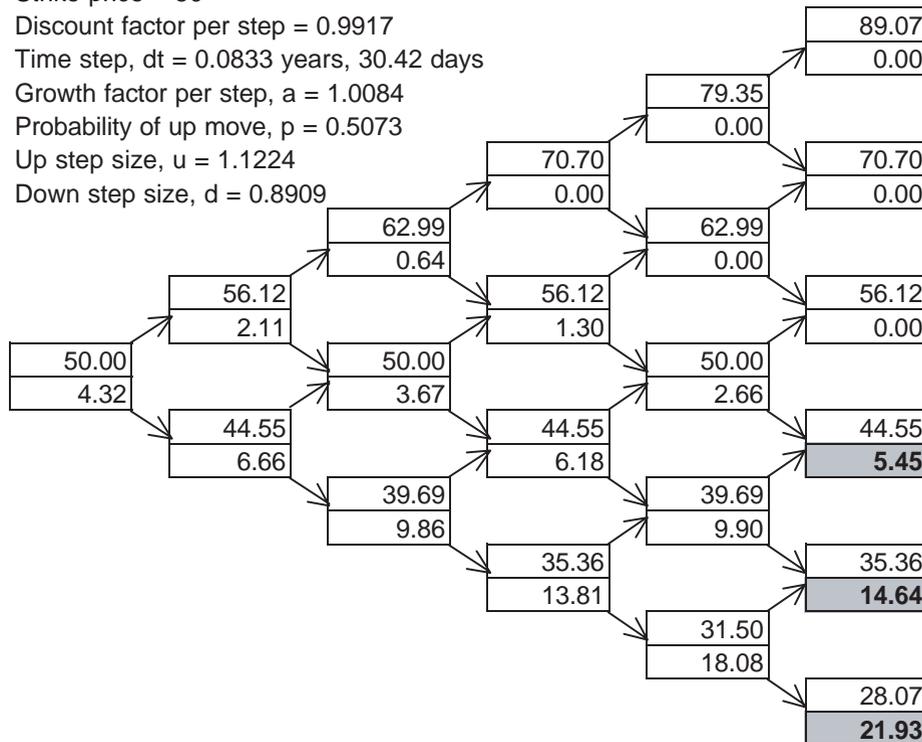
Time step, dt = 0.0833 years, 30.42 days

Growth factor per step, a = 1.0084

Probability of up move, p = 0.5073

Up step size, u = 1.1224

Down step size, d = 0.8909



Node Time:

0.0000

0.0833

0.1667

0.2500

0.3333

0.4167

20.4 ALTERNATIVE PROCEDURES FOR CONSTRUCTING TREES

The Cox, Ross, and Rubinstein approach is not the only way of building a binomial tree. Instead of imposing the assumption $u = 1/d$ on equations (20.1) and (20.2), we can set $p = 0.5$. A solution to the equations when terms of higher order than Δt are ignored is then

$$u = e^{(r-q-\sigma^2/2)\Delta t + \sigma\sqrt{\Delta t}}, \quad d = e^{(r-q-\sigma^2/2)\Delta t - \sigma\sqrt{\Delta t}}$$

This allows trees with $p = 0.5$ to be built for options on stocks, indices, foreign exchange, and futures.

This alternative tree-building procedure has the advantage over the Cox, Ross, and Rubinstein approach that the probabilities are always 0.5 regardless of the value of σ or the number of time steps.⁹ Its disadvantage is that it is not quite as straightforward to calculate delta, gamma, and rho from the tree because the tree is no longer centered at the initial stock price.

Example 20.6

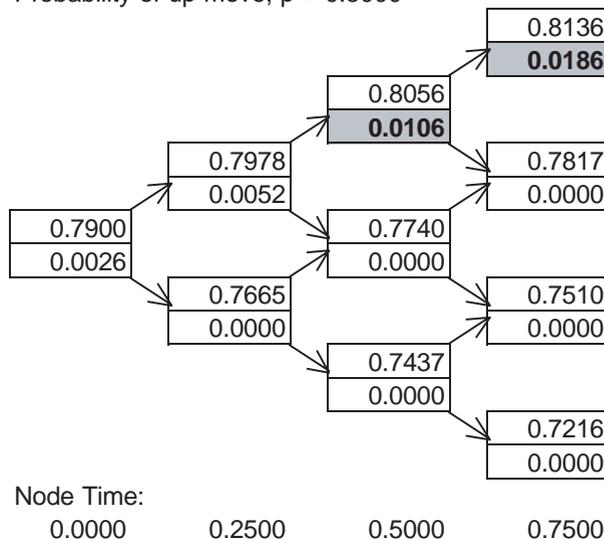
Consider a 9-month American call option on a foreign currency. The foreign currency is worth 0.7900 when measured in the domestic currency, the strike price is 0.7950, the domestic risk-free interest rate is 6% per annum, the foreign risk-free interest rate is 10% per annum, and the volatility of the exchange rate is 4% per annum. In this case, $S_0 = 0.79$, $K = 0.795$, $r = 0.06$, $r_f = 0.10$, $\sigma = 0.04$, and

Figure 20.11 Binomial tree for American call option on a foreign currency. At each node, upper number is spot exchange rate and lower number is option price. All probabilities are 0.5.

At each node:
 Upper value = Underlying Asset Price
 Lower value = Option Price
 Shading indicates where option is exercised

Strike price = 0.795
 Discount factor per step = 0.9851
 Time step, $dt = 0.2500$ years, 91.25 days

Probability of up move, $p = 0.5000$



⁹ When time steps are so large that $\sigma < |(r - q)\sqrt{\Delta t}|$, the Cox, Ross, and Rubinstein tree gives negative probabilities. The alternative procedure described here does not have that drawback.

$T = 0.75$. We set $\Delta t = 0.25$ (3 steps) and the probabilities on each branch to 0.5, so that

$$u = e^{(0.06 - 0.10 - 0.0016/2)0.25 + 0.04\sqrt{0.25}} = 1.0098$$

$$d = e^{(0.06 - 0.10 - 0.0016/2)0.25 - 0.04\sqrt{0.25}} = 0.9703$$

The tree for the exchange rate is shown in Figure 20.11. The tree gives the value of the option as \$0.0026.

Trinomial Trees

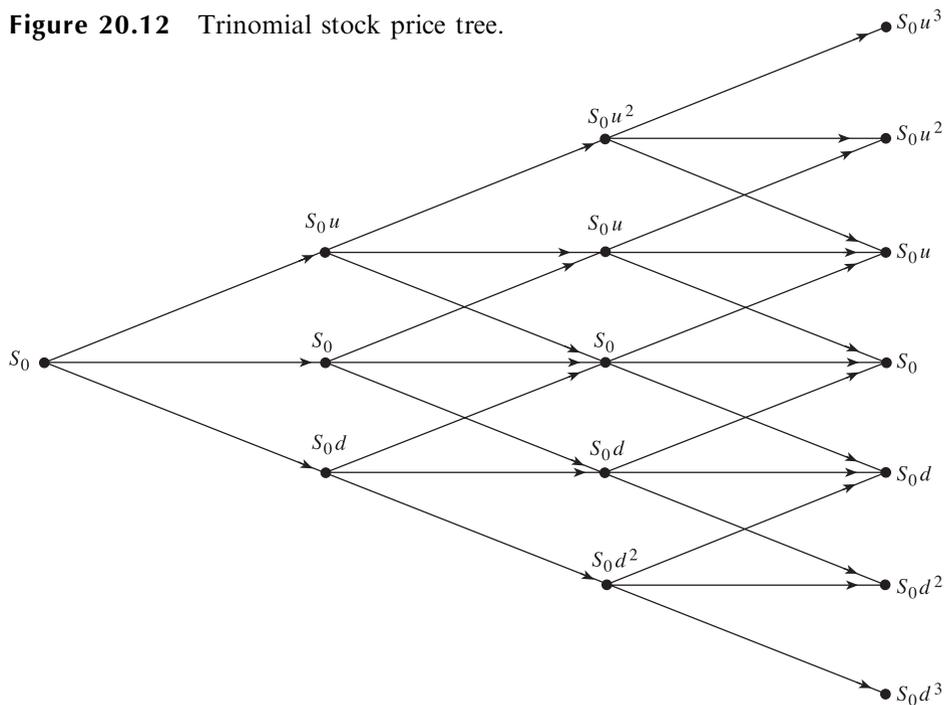
Trinomial trees can be used as an alternative to binomial trees. The general form of the tree is as shown in Figure 20.12. Suppose that p_u , p_m , and p_d are the probabilities of up, middle, and down movements at each node and Δt is the length of the time step. For an asset paying dividends at a rate q , parameter values that match the mean and standard deviation of price changes when terms of higher order than Δt are ignored are

$$u = e^{\sigma\sqrt{3\Delta t}}, \quad d = 1/u$$

$$p_d = -\sqrt{\frac{\Delta t}{12\sigma^2}}\left(r - q - \frac{\sigma^2}{2}\right) + \frac{1}{6}, \quad p_m = \frac{2}{3}, \quad p_u = \sqrt{\frac{\Delta t}{12\sigma^2}}\left(r - q - \frac{\sigma^2}{2}\right) + \frac{1}{6}$$

Calculations for a trinomial tree are analogous to those for a binomial tree. We work from the end of the tree to the beginning. At each node we calculate the value of

Figure 20.12 Trinomial stock price tree.



exercising and the value of continuing. The value of continuing is

$$e^{-r\Delta t}(p_u f_u + p_m f_m + p_d f_d)$$

where f_u , f_m , and f_d are the values of the option at the subsequent up, middle, and down nodes, respectively. The trinomial tree approach proves to be equivalent to the explicit finite difference method, which will be described in Section 20.8.

Figlewski and Gao have proposed an enhancement of the trinomial tree method, which they call the *adaptive mesh model*. In this, a high-resolution (small- Δt) tree is grafted onto a low-resolution (large- Δt) tree.¹⁰ When valuing a regular American option, high resolution is most useful for the parts of the tree close to the strike price at the end of the life of the option.

20.5 TIME-DEPENDENT PARAMETERS

Up to now we have assumed that r , q , r_f , and σ are constants. In practice, they are usually assumed to be time dependent. The values of these variables between times t and $t + \Delta t$ are assumed to be equal to their forward values.¹¹

To make r and q (or r_f) a function of time in a Cox–Ross–Rubinstein binomial tree, we set

$$a = e^{[f(t)-g(t)]\Delta t} \quad (20.11)$$

for nodes at time t , where $f(t)$ is the forward interest rate between times t and $t + \Delta t$ and $g(t)$ is the forward value of q (or r_f) between these times. This does not change the geometry of the tree because u and d do not depend on a . The probabilities on the branches emanating from nodes at time t are:¹²

$$p = \frac{e^{[f(t)-g(t)]\Delta t} - d}{u - d} \quad (20.12)$$

$$1 - p = \frac{u - e^{[f(t)-g(t)]\Delta t}}{u - d}$$

The rest of the way that we use the tree is the same as before, except that when discounting between times t and $t + \Delta t$ we use $f(t)$.

Making σ a function of time in a binomial tree is more challenging. One approach is to make the lengths of time steps inversely proportional to the variance rate. The values of u and d are then always the same and the tree recombines. Suppose that $\sigma(t)$ is the volatility for a maturity t so that $\sigma(t)^2 t$ is the cumulative variance by time t . Define $V = \sigma(T)^2 T$, where T is the life of the tree, and let t_i be the end of the i th time step. If there is a total of N time steps, we choose t_i to satisfy $\sigma(t_i)^2 t_i = iV/N$. The variance between times t_{i-1} and t_i is then V/N for all i .

¹⁰ See S. Figlewski and B. Gao, “The Adaptive Mesh Model: A New Approach to Efficient Option Pricing,” *Journal of Financial Economics*, 53 (1999): 313–51.

¹¹ The forward dividend yield and forward variance rate are calculated in the same way as the forward interest rate. (The variance rate is the square of the volatility.)

¹² For a sufficiently large number of time steps, these probabilities are always positive.

Business Snapshot 20.1 Calculating Pi with Monte Carlo Simulation

Suppose the sides of the square in Figure 20.13 are one unit in length. Imagine that you fire darts randomly at the square and calculate the percentage that lie in the circle. What should you find? The square has an area of 1.0 and the circle has a radius of 0.5. The area of the circle is π times the radius squared or $\pi/4$. It follows that the proportion of darts that lie in the circle should be $\pi/4$. We can estimate π by multiplying the proportion that lie in the circle by 4.

We can use an Excel spreadsheet to simulate the dart throwing as illustrated in Table 20.1. We define both cell A1 and cell B1 as `=RAND()`. A1 and B1 are random numbers between 0 and 1 and define how far to the right and how high up the dart lands in the square in Figure 20.13. We then define cell C1 as

$$=IF((A1-0.5)^2+(B1-0.5)^2<0.5^2,4,0)$$

This has the effect of setting C1 equal to 4 if the dart lies in the circle and 0 otherwise.

Define the next 99 rows of the spreadsheet similarly to the first one. (This is a “select and drag” operation in Excel.) Define C102 as `=AVERAGE(C1:C100)` and C103 as `=STDEV(C1:C100)`. C102 (which is 3.04 in Table 20.1) is an estimate of π calculated from 100 random trials. C103 is the standard deviation of our results and as we will see in Example 20.7 can be used to assess the accuracy of the estimate. Increasing the number of trials improves accuracy—but convergence to the correct value of 3.14159 is slow.

With a trinomial tree, a generalized tree-building procedure can be used to match time-dependent interest rates and volatilities (see Technical Note 9 on the author’s website).

20.6 MONTE CARLO SIMULATION

We now explain Monte Carlo simulation, a quite different approach for valuing derivatives from binomial trees. Business Snapshot 20.1 illustrates the random sampling idea underlying Monte Carlo simulation by showing how a simple Excel program can be constructed to estimate π .

When used to value an option, Monte Carlo simulation uses the risk-neutral valuation result. We sample paths to obtain the expected payoff in a risk-neutral world

Figure 20.13 Calculation of π by throwing darts.

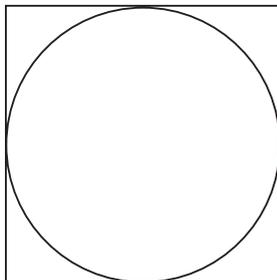


Table 20.1 Sample spreadsheet calculations in Business Snapshot 20.1.

	<i>A</i>	<i>B</i>	<i>C</i>
1	0.207	0.690	4
2	0.271	0.520	4
3	0.007	0.221	0
⋮	⋮	⋮	⋮
100	0.198	0.403	4
101			
102		Mean:	3.04
103		SD:	1.69

and then discount this payoff at the risk-free rate. Consider a derivative dependent on a single market variable S that provides a payoff at time T . Assuming that interest rates are constant, we can value the derivative as follows:

1. Sample a random path for S in a risk-neutral world.
2. Calculate the payoff from the derivative.
3. Repeat steps 1 and 2 to get many sample values of the payoff from the derivative in a risk-neutral world.
4. Calculate the mean of the sample payoffs to get an estimate of the expected payoff in a risk-neutral world.
5. Discount this expected payoff at the risk-free rate to get an estimate of the value of the derivative.

Suppose that the process followed by the underlying market variable in a risk-neutral world is

$$dS = \hat{\mu}S dt + \sigma S dz \quad (20.13)$$

where dz is a Wiener process, $\hat{\mu}$ is the expected return in a risk-neutral world, and σ is the volatility.¹³ To simulate the path followed by S , we can divide the life of the derivative into N short intervals of length Δt and approximate equation (20.13) as

$$S(t + \Delta t) - S(t) = \hat{\mu}S(t) \Delta t + \sigma S(t)\epsilon\sqrt{\Delta t} \quad (20.14)$$

where $S(t)$ denotes the value of S at time t , ϵ is a random sample from a normal distribution with mean zero and standard deviation of 1.0. This enables the value of S at time Δt to be calculated from the initial value of S , the value at time $2\Delta t$ to be calculated from the value at time Δt , and so on. An illustration of the procedure is in Section 13.3. One simulation trial involves constructing a complete path for S using N random samples from a normal distribution.

¹³ If S is the price of a non-dividend-paying stock then $\hat{\mu} = r$, if it is an exchange rate then $\hat{\mu} = r - r_f$, and so on. Note that the volatility is the same in a risk-neutral world as in the real world, as explained in Section 12.7.

In practice, it is usually more accurate to simulate $\ln S$ rather than S . From Itô's lemma the process followed by $\ln S$ is

$$d \ln S = \left(\hat{\mu} - \frac{\sigma^2}{2} \right) dt + \sigma dz \quad (20.15)$$

so that

$$\ln S(t + \Delta t) - \ln S(t) = \left(\hat{\mu} - \frac{\sigma^2}{2} \right) \Delta t + \sigma \epsilon \sqrt{\Delta t}$$

or equivalently

$$S(t + \Delta t) = S(t) \exp \left[\left(\hat{\mu} - \frac{\sigma^2}{2} \right) \Delta t + \sigma \epsilon \sqrt{\Delta t} \right] \quad (20.16)$$

This equation is used to construct a path for S .

Working with $\ln S$ rather than S gives more accuracy. Also, if $\hat{\mu}$ and σ are constant, then

$$\ln S(T) - \ln S(0) = \left(\hat{\mu} - \frac{\sigma^2}{2} \right) T + \sigma \epsilon \sqrt{T}$$

is true for all T .¹⁴ It follows that

$$S(T) = S(0) \exp \left[\left(\hat{\mu} - \frac{\sigma^2}{2} \right) T + \sigma \epsilon \sqrt{T} \right] \quad (20.17)$$

This equation can be used to value derivatives that provide a nonstandard payoff at time T . As shown in Business Snapshot 20.2, it can also be used to check the Black–Scholes–Merton formulas.

The key advantage of Monte Carlo simulation is that it can be used when the payoff depends on the path followed by the underlying variable S as well as when it depends only on the final value of S . (For example, it can be used when payoffs depend on the average value of S between time 0 and time T .) Payoffs can occur at several times during the life of the derivative rather than all at the end. Any stochastic process for S can be accommodated. As will be shown shortly, the procedure can also be extended to accommodate situations where the payoff from the derivative depends on several underlying market variables. The drawbacks of Monte Carlo simulation are that it is computationally very time consuming and cannot easily handle situations where there are early exercise opportunities.¹⁵

Derivatives Dependent on More than One Market Variable

We discussed correlated stochastic processes in Section 13.5. Consider the situation where the payoff from a derivative depends on n variables θ_i ($1 \leq i \leq n$). Define s_i as the volatility of θ_i , \hat{m}_i as the expected growth rate of θ_i in a risk-neutral world, and ρ_{ik} as the correlation between the Wiener processes driving θ_i and θ_k .¹⁶ As in the single-variable case, the life of the derivative must be divided into N subintervals of length Δt . The

¹⁴ By contrast, equation (20.14) is exactly true only in the limit as Δt tends to zero.

¹⁵ As discussed in Chapter 26, a number of researchers have suggested ways Monte Carlo simulation can be extended to value American options.

¹⁶ Note that s_i , \hat{m}_i , and ρ_{ik} are not necessarily constant; they may depend on the θ_j .

Business Snapshot 20.2 Checking Black–Scholes–Merton in Excel

The Black–Scholes–Merton formula for a European call option can be checked by using a binomial tree with a very large number of time steps. An alternative way of checking it is to use Monte Carlo simulation. Table 20.2 shows a spreadsheet that can be constructed. The cells C2, D2, E2, F2, and G2 contain S_0 , K , r , σ , and T , respectively. Cells D4, E4, and F4 calculate d_1 , d_2 , and the Black–Scholes–Merton price, respectively. (The Black–Scholes–Merton price is 4.817 in the sample spreadsheet.)

NORMSINV is the inverse cumulative function for the standard normal distribution. It follows that NORMSINV(RAND()) gives a random sample from a standard normal distribution. We set cell A1 as

$$= \$C\$2 * EXP((\$E\$2 - \$F\$2 * \$F\$2 / 2) * \$G\$2 + \$F\$2 * NORMSINV(RAND()) * SQRT(\$G\$2))$$

This corresponds to equation (20.17) and is a random sample from the set of all stock prices at time T . We set cell B1 as

$$= EXP(-\$E\$2 * \$G\$2) * MAX(A1 - \$D\$2, 0)$$

This is the present value of the payoff from a call option. We define the next 999 rows of the spreadsheet similarly to the first one. (This is a “select and drag” operation in Excel.) Define B1002 as AVERAGE(B1:B1000) and B1003 as STDEV(B1:B1000). B1002 (which is 4.98 in the sample spreadsheet) is an estimate of the value of the option. This should be not too far from the Black–Scholes–Merton price. As we shall see in Example 20.8, B1003 can be used to assess the accuracy of the estimate.

discrete version of the process for θ_i is then

$$\theta_i(t + \Delta t) - \theta_i(t) = \hat{m}_i \theta_i(t) \Delta t + s_i \theta_i(t) \epsilon_i \sqrt{\Delta t} \tag{20.18}$$

where ϵ_i is a random sample from a standard normal distribution. The coefficient of correlation between ϵ_i and ϵ_k is ρ_{ik} ($1 \leq i, k \leq n$). One simulation trial involves obtaining N samples of the ϵ_i ($1 \leq i \leq n$) from a multivariate standardized normal distribution. These are substituted into equation (20.18) to produce simulated paths for each θ_i , thereby enabling a sample value for the derivative to be calculated.

Table 20.2 Monte Carlo simulation to check Black–Scholes–Merton.

	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	<i>E</i>	<i>F</i>	<i>G</i>
1	45.95	0	S_0	K	r	σ	T
2	54.49	4.38	50	50	0.05	0.3	0.5
3	50.09	0.09		d_1	d_2	BSM price	
4	47.46	0		0.2239	0.0118	4.817	
5	44.93	0					
⋮	⋮	⋮					
1000	68.27	17.82					
1001							
1002	Mean:	4.98					
1003	SD:	7.68					

Generating the Random Samples from Normal Distributions

The instruction =NORMSINV(RAND()) in Excel can be used to generate a random sample from a standard normal distribution, as in Business Snapshot 20.2. When two correlated samples ϵ_1 and ϵ_2 from standard normal distributions are required, an appropriate procedure is as follows. Independent samples x_1 and x_2 from a univariate standardized normal distribution are obtained as just described. The required samples ϵ_1 and ϵ_2 are then calculated as follows:

$$\begin{aligned}\epsilon_1 &= x_1 \\ \epsilon_2 &= \rho x_1 + x_2 \sqrt{1 - \rho^2}\end{aligned}$$

where ρ is the coefficient of correlation.

More generally, consider the situation where we require n correlated samples from normal distributions with the correlation between sample i and sample j being ρ_{ij} . We first sample n independent variables x_i ($1 \leq i \leq n$), from univariate standardized normal distributions. The required samples, ϵ_i ($1 \leq i \leq n$), are then defined as follows:

$$\left. \begin{aligned}\epsilon_1 &= \alpha_{11}x_1 \\ \epsilon_2 &= \alpha_{21}x_1 + \alpha_{22}x_2 \\ \epsilon_3 &= \alpha_{31}x_1 + \alpha_{32}x_2 + \alpha_{33}x_3\end{aligned}\right\} \quad (20.19)$$

and so on. We choose the coefficients α_{ij} so that the correlations and variances are correct. This can be done step by step as follows. Set $\alpha_{11} = 1$; choose α_{21} so that $\alpha_{21}\alpha_{11} = \rho_{21}$; choose α_{22} so that $\alpha_{21}^2 + \alpha_{22}^2 = 1$; choose α_{31} so that $\alpha_{31}\alpha_{11} = \rho_{31}$; choose α_{32} so that $\alpha_{31}\alpha_{21} + \alpha_{32}\alpha_{22} = \rho_{32}$; choose α_{33} so that $\alpha_{31}^2 + \alpha_{32}^2 + \alpha_{33}^2 = 1$; and so on.¹⁷ This procedure is known as the *Cholesky decomposition*.

Number of Trials

The accuracy of the result given by Monte Carlo simulation depends on the number of trials. It is usual to calculate the standard deviation as well as the mean of the discounted payoffs given by the simulation trials. Denote the mean by μ and the standard deviation by ω . The variable μ is the simulation's estimate of the value of the derivative. The standard error of the estimate is

$$\frac{\omega}{\sqrt{M}}$$

where M is the number of trials. A 95% confidence interval for the price f of the derivative is therefore given by

$$\mu - \frac{1.96\omega}{\sqrt{M}} < f < \mu + \frac{1.96\omega}{\sqrt{M}}$$

This shows that uncertainty about the value of the derivative is inversely proportional to the square root of the number of trials. To double the accuracy of a simulation, we

¹⁷ If the equations for the α 's do not have real solutions, the assumed correlation structure is internally inconsistent. This will be discussed further in Section 22.7.

must quadruple the number of trials; to increase the accuracy by a factor of 10, the number of trials must increase by a factor of 100; and so on.

Example 20.7

In Table 20.1, π is calculated as the average of 100 numbers. The standard deviation of the numbers is 1.69. In this case, $\omega = 1.69$ and $M = 100$, so that the standard error of the estimate is $1.69/\sqrt{100} = 0.169$. The spreadsheet therefore gives a 95% confidence interval for π as $(3.04 - 1.96 \times 0.169)$ to $(3.04 + 1.96 \times 0.169)$ or 2.71 to 3.37. (The correct value of 3.14159 lies within this confidence interval.)

Example 20.8

In Table 20.2, the value of the option is calculated as the average of 1000 numbers. The standard deviation of the numbers is 7.68. In this case, $\omega = 7.68$ and $M = 1000$. The standard error of the estimate is $7.68/\sqrt{1000} = 0.24$. The spreadsheet therefore gives a 95% confidence interval for the option value as $(4.98 - 1.96 \times 0.24)$ to $(4.98 + 1.96 \times 0.24)$, or 4.51 to 5.45. (The Black–Scholes–Merton price, 4.817, lies within this confidence interval.)

Sampling through a Tree

Instead of implementing Monte Carlo simulation by randomly sampling from the stochastic process for an underlying variable, we can use an N -step binomial tree and sample from the 2^N paths that are possible. Suppose we have a binomial tree where the probability of an “up” movement is 0.6. The procedure for sampling a random path through the tree is as follows. At each node, we sample a random number between 0 and 1. If the number is less than 0.4, we take the down branch. If it is greater than 0.4, we take the up branch. Once we have a complete path from the initial node to the end of the tree, we can calculate a payoff. This completes the first trial. A similar procedure is used to complete more trials. The mean of the payoffs is discounted at the risk-free rate to get an estimate of the value of the derivative.¹⁸

Example 20.9

Suppose that the tree in Figure 20.3 is used to value an option that pays off $\max(S_{\text{ave}} - 50, 0)$, where S_{ave} is the average stock price during the 5 months (with the first and last stock price being included in the average). This is known as an Asian option. When ten simulation trials are used one possible result is shown in Table 20.3. The value of the option is calculated as the average payoff discounted at the risk-free rate. In this case, the average payoff is \$7.08 and the risk-free rate is 10% and so the calculated value is $7.08e^{-0.1 \times 5/12} = 6.79$. (This illustrates the methodology. In practice, we would have to use more time steps on the tree and many more simulation trials to get an accurate answer.)

Calculating the Greek Letters

The Greek letters discussed in Chapter 18 can be calculated using Monte Carlo simulation. Suppose that we are interested in the partial derivative of f with respect

¹⁸ See D. Mintz, “Less is More,” *Risk*, July 1997: 42–45, for a discussion of how sampling through a tree can be made efficient.

Table 20.3 Monte Carlo simulation to value Asian option from the tree in Figure 19.3. Payoff is amount by which average stock price exceeds \$50. U = up movement; D = down movement.

<i>Trial</i>	<i>Path</i>	<i>Average stock price</i>	<i>Option payoff</i>
1	UUUUD	64.98	14.98
2	UUUDD	59.82	9.82
3	DDDUU	42.31	0.00
4	UUUUU	68.04	18.04
5	UUDDU	55.22	5.22
6	UDUUD	55.22	5.22
7	DDUDD	42.31	0.00
8	UUDDU	55.22	5.22
9	UUUDU	62.25	12.25
10	DDUUD	45.56	0.00
Average			7.08

to x , where f is the value of the derivative and x is the value of an underlying variable or a parameter. First, Monte Carlo simulation is used in the usual way to calculate an estimate \hat{f} for the value of the derivative. A small increase Δx is then made in the value of x , and a new value for the derivative, \hat{f}^* , is calculated in the same way as \hat{f} . An estimate for the hedge parameter is given by

$$\frac{\hat{f}^* - \hat{f}}{\Delta x}$$

In order to minimize the standard error of the estimate, the number of time intervals, N , the random samples that are used, and the number of trials, M , should be the same for calculating both \hat{f} and \hat{f}^* .

Applications

Monte Carlo simulation tends to be numerically more efficient than other procedures when there are three or more stochastic variables. This is because the time taken to carry out a Monte Carlo simulation increases approximately linearly with the number of variables, whereas the time taken for most other procedures increases exponentially with the number of variables. One advantage of Monte Carlo simulation is that it can provide a standard error for the estimates that it makes. Another is that it is an approach that can accommodate complex payoffs and complex stochastic processes. Also, it can be used when the payoff depends on some function of the whole path followed by a variable, not just its terminal value.

20.7 VARIANCE REDUCTION PROCEDURES

If the stochastic processes for the variables underlying a derivative are simulated as indicated in equations (20.13) to (20.18), a very large number of trials is usually

necessary to estimate the value of the derivative with reasonable accuracy. This is very expensive in terms of computation time. In this section, we examine a number of variance reduction procedures that can lead to dramatic savings in computation time.

Antithetic Variable Technique

In the antithetic variable technique, a simulation trial involves calculating two values of the derivative. The first value f_1 is calculated in the usual way; the second value f_2 is calculated by changing the sign of all the random samples from standard normal distributions. (If ϵ is a sample used to calculate f_1 , then $-\epsilon$ is the corresponding sample used to calculate f_2 .) The sample value of the derivative calculated from a simulation trial is the average of f_1 and f_2 . This works well because when one value is above the true value, the other tends to be below, and vice versa.

Denote \bar{f} as the average of f_1 and f_2 :

$$\bar{f} = \frac{f_1 + f_2}{2}$$

The final estimate of the value of the derivative is the average of the \bar{f} 's. If $\bar{\omega}$ is the standard deviation of the \bar{f} 's, and M is the number of simulation trials (i.e., the number of pairs of values calculated), then the standard error of the estimate is

$$\bar{\omega}/\sqrt{M}$$

This is usually much less than the standard error calculated using $2M$ random trials.

Control Variate Technique

We have already given one example of the control variate technique in connection with the use of trees to value American options (see Section 20.3). The control variate technique is applicable when there are two similar derivatives, A and B. Derivative A is the one being valued; derivative B is similar to derivative A and has an analytic solution available. Two simulations using the same random number streams and the same Δt are carried out in parallel. The first is used to obtain an estimate f_A^* of the value of A; the second is used to obtain an estimate f_B^* , of the value of B. A better estimate f_A of the value of A is then obtained using the formula

$$f_A = f_A^* - f_B^* + f_B \quad (20.20)$$

where f_B is the known true value of B calculated analytically. Hull and White provide an example of the use of the control variate technique when evaluating the effect of stochastic volatility on the price of a European call option.¹⁹ In this case, A is the option assuming stochastic volatility and B is the option assuming constant volatility.

Importance Sampling

Importance sampling is best explained with an example. Suppose that we wish to calculate the price of a deep-out-of-the-money European call option with strike

¹⁹ See J. Hull and A. White, "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance*, 42 (June 1987): 281–300.

price K and maturity T . If we sample values for the underlying asset price at time T in the usual way, most of the paths will lead to zero payoff. This is a waste of computation time because the zero-payoff paths contribute very little to the determination of the value of the option. We therefore try to choose only important paths, that is, paths where the stock price is above K at maturity.

Suppose F is the unconditional probability distribution function for the stock price at time T and q , the probability of the stock price being greater than K at maturity, is known analytically. Then $G = F/q$ is the probability distribution of the stock price conditional on the stock price being greater than K . To implement importance sampling, we sample from G rather than F . The estimate of the value of the option is the average discounted payoff multiplied by q .

Stratified Sampling

Sampling representative values rather than random values from a probability distribution usually gives more accuracy. Stratified sampling is a way of doing this. Suppose we wish to take 1000 samples from a probability distribution. We would divide the distribution into 1000 equally likely intervals and choose a representative value (typically the mean or median) for each interval.

In the case of a standard normal distribution when there are n intervals, we can calculate the representative value for the i th interval as

$$N^{-1}\left(\frac{i-0.5}{n}\right)$$

where N^{-1} is the inverse cumulative normal distribution. For example, when $n = 4$ the representative values corresponding to the four intervals are $N^{-1}(0.125)$, $N^{-1}(0.375)$, $N^{-1}(0.625)$, $N^{-1}(0.875)$. The function N^{-1} can be calculated using the NORMSINV function in Excel.

Moment Matching

Moment matching involves adjusting the samples taken from a standardized normal distribution so that the first, second, and possibly higher moments are matched. Suppose that we sample from a normal distribution with mean 0 and standard deviation 1 to calculate the change in the value of a particular variable over a particular time period. Suppose that the samples are ϵ_i ($1 \leq i \leq n$). To match the first two moments, we calculate the mean of the samples, m , and the standard deviation of the samples, s . We then define adjusted samples ϵ_i^* ($1 \leq i \leq n$) as

$$\epsilon_i^* = \frac{\epsilon_i - m}{s}$$

These adjusted samples have the correct mean of 0 and the correct standard deviation of 1.0. We use the adjusted samples for all calculations.

Moment matching saves computation time, but can lead to memory problems because every number sampled must be stored until the end of the simulation. Moment matching is sometimes termed *quadratic resampling*. It is often used in conjunction with the antithetic variable technique. Because the latter automatically matches all odd

moments, the goal of moment matching then becomes that of matching the second moment and, possibly, the fourth moment.

Using Quasi-Random Sequences

A quasi-random sequence (also called a *low-discrepancy* sequence) is a sequence of representative samples from a probability distribution.²⁰ Descriptions of the use of quasi-random sequences appear in Brotherton-Ratcliffe, and Press *et al.*²¹ Quasi-random sequences can have the desirable property that they lead to the standard error of an estimate being proportional to $1/M$ rather than $1/\sqrt{M}$, where M is the sample size.

Quasi-random sampling is similar to stratified sampling. The objective is to sample representative values for the underlying variables. In stratified sampling, it is assumed that we know in advance how many samples will be taken. A quasi-random sampling procedure is more flexible. The samples are taken in such a way that we are always “filling in” gaps between existing samples. At each stage of the simulation, the sampled points are roughly evenly spaced throughout the probability space.

Figure 20.14 shows points generated in two dimensions using a procedure suggested by Sobol’.²² It can be seen that successive points do tend to fill in the gaps left by previous points.

20.8 FINITE DIFFERENCE METHODS

Finite difference methods value a derivative by solving the differential equation that the derivative satisfies. The differential equation is converted into a set of difference equations, and the difference equations are solved iteratively.

To illustrate the approach, we consider how it might be used to value an American put option on a stock paying a dividend yield of q . The differential equation that the option must satisfy is, from equation (16.6),

$$\frac{\partial f}{\partial t} + (r - q)S \frac{\partial f}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} = rf \quad (20.21)$$

Suppose that the life of the option is T . We divide this into N equally spaced intervals of length $\Delta t = T/N$. A total of $N + 1$ times are therefore considered

$$0, \Delta t, 2\Delta t, \dots, T$$

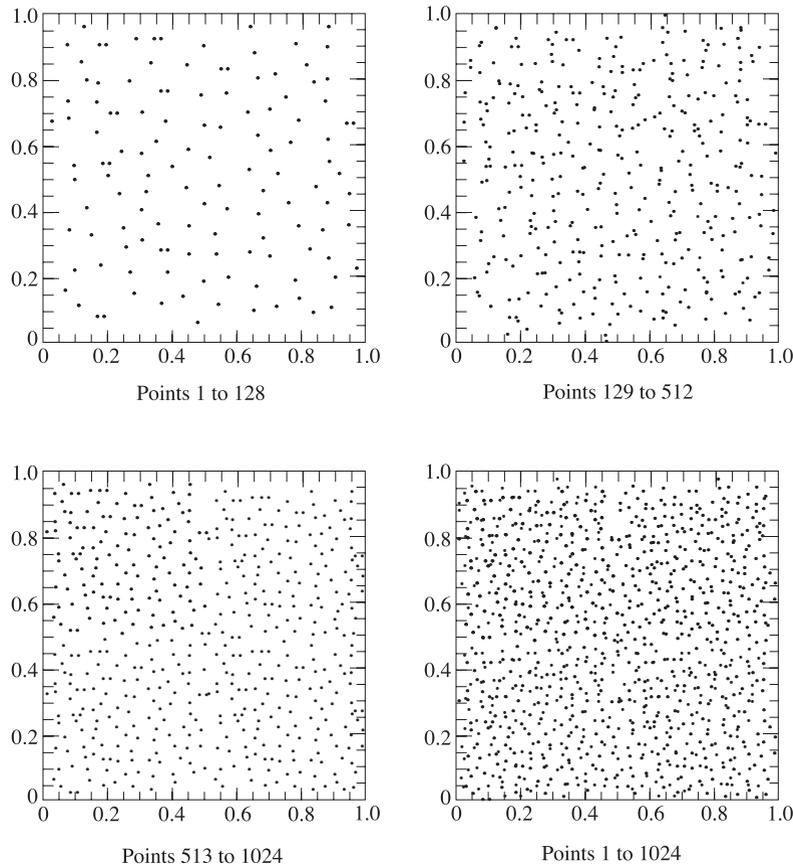
Suppose that S_{\max} is a stock price sufficiently high that, when it is reached, the put has virtually no value. We define $\Delta S = S_{\max}/M$ and consider a total of $M + 1$ equally spaced stock prices:

$$0, \Delta S, 2\Delta S, \dots, S_{\max}$$

²⁰ The term *quasi-random* is a misnomer. A quasi-random sequence is totally deterministic.

²¹ See R. Brotherton-Ratcliffe, “Monte Carlo Motoring,” *Risk*, December 1994: 53–58; W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*, 2nd edn. Cambridge University Press, 1992.

²² See I. M. Sobol’, *USSR Computational Mathematics and Mathematical Physics*, 7, 4 (1967): 86–112. A description of Sobol’'s procedure is in W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*, 2nd edn. Cambridge University Press, 1992.

Figure 20.14 First 1024 points of a Sobol' sequence.

The level S_{\max} is chosen so that one of these is the current stock price.

The time points and stock price points define a grid consisting of a total of $(M + 1)(N + 1)$ points, as shown in Figure 20.15. The (i, j) point on the grid is the point that corresponds to time $i \Delta t$ and stock price $j \Delta S$. We will use the variable $f_{i,j}$ to denote the value of the option at the (i, j) point.

Implicit Finite Difference Method

For an interior point (i, j) on the grid, $\partial f / \partial S$ can be approximated as

$$\frac{\partial f}{\partial S} = \frac{f_{i,j+1} - f_{i,j}}{\Delta S} \quad (20.22)$$

or as

$$\frac{\partial f}{\partial S} = \frac{f_{i,j} - f_{i,j-1}}{\Delta S} \quad (20.23)$$

Equation (20.22) is known as the *forward difference approximation*; equation (20.23) is known as the *backward difference approximation*. We use a more symmetrical

approximation by averaging the two:

$$\frac{\partial f}{\partial S} = \frac{f_{i,j+1} - f_{i,j-1}}{2 \Delta S} \tag{20.24}$$

For $\partial f / \partial t$, we will use a forward difference approximation so that the value at time $i \Delta t$ is related to the value at time $(i + 1) \Delta t$:

$$\frac{\partial f}{\partial t} = \frac{f_{i+1,j} - f_{i,j}}{\Delta t} \tag{20.25}$$

Consider next $\partial^2 f / \partial S^2$. The backward difference approximation for $\partial f / \partial S$ at the (i, j) point is given by equation (20.23). The backward difference at the $(i, j + 1)$ point is

$$\frac{f_{i,j+1} - f_{i,j}}{\Delta S}$$

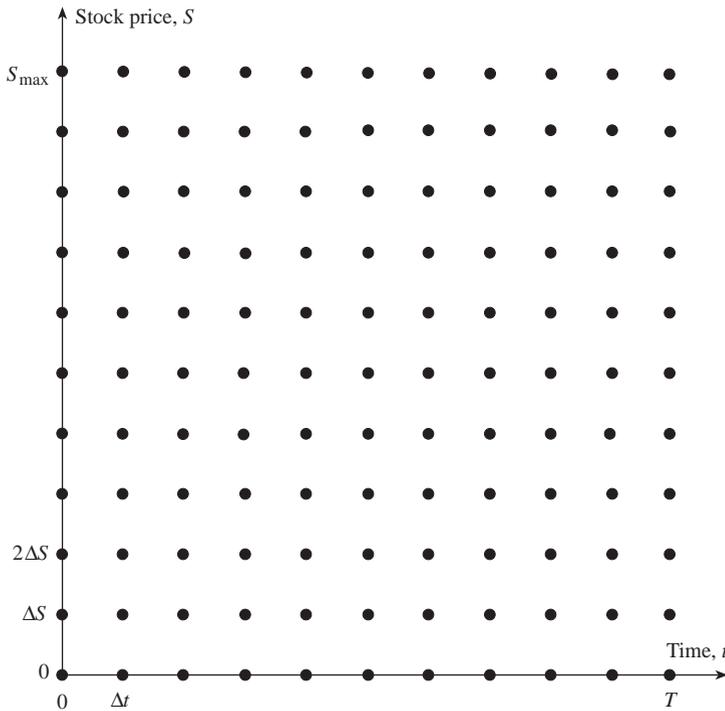
Hence a finite difference approximation for $\partial^2 f / \partial S^2$ at the (i, j) point is

$$\frac{\partial^2 f}{\partial S^2} = \left(\frac{f_{i,j+1} - f_{i,j}}{\Delta S} - \frac{f_{i,j} - f_{i,j-1}}{\Delta S} \right) / \Delta S$$

or

$$\frac{\partial^2 f}{\partial S^2} = \frac{f_{i,j+1} + f_{i,j-1} - 2f_{i,j}}{\Delta S^2} \tag{20.26}$$

Figure 20.15 Grid for finite difference approach.



Substituting equations (20.24), (20.25), and (20.26) into the differential equation (20.21) and noting that $S = j \Delta S$ gives

$$\frac{f_{i+1,j} - f_{i,j}}{\Delta t} + (r - q)j \Delta S \frac{f_{i,j+1} - f_{i,j-1}}{2 \Delta S} + \frac{1}{2} \sigma^2 j^2 \Delta S^2 \frac{f_{i,j+1} + f_{i,j-1} - 2f_{i,j}}{\Delta S^2} = r f_{i,j}$$

for $j = 1, 2, \dots, M - 1$ and $i = 0, 1, \dots, N - 1$. Rearranging terms, we obtain

$$a_j f_{i,j-1} + b_j f_{i,j} + c_j f_{i,j+1} = f_{i+1,j} \quad (20.27)$$

where

$$a_j = \frac{1}{2}(r - q)j \Delta t - \frac{1}{2} \sigma^2 j^2 \Delta t$$

$$b_j = 1 + \sigma^2 j^2 \Delta t + r \Delta t$$

$$c_j = -\frac{1}{2}(r - q)j \Delta t - \frac{1}{2} \sigma^2 j^2 \Delta t$$

The value of the put at time T is $\max(K - S_T, 0)$, where S_T is the stock price at time T . Hence,

$$f_{N,j} = \max(K - j \Delta S, 0), \quad j = 0, 1, \dots, M \quad (20.28)$$

The value of the put option when the stock price is zero is K . Hence,

$$f_{i,0} = K, \quad i = 0, 1, \dots, N \quad (20.29)$$

We assume that the put option is worth zero when $S = S_{\max}$, so that

$$f_{i,M} = 0, \quad i = 0, 1, \dots, N \quad (20.30)$$

Equations (20.28), (20.29), and (20.30) define the value of the put option along the three edges of the grid in Figure 20.15, where $S = 0$, $S = S_{\max}$, and $t = T$. It remains to use equation (20.27) to arrive at the value of f at all other points. First the points corresponding to time $T - \Delta t$ are tackled. Equation (20.27) with $i = N - 1$ gives

$$a_j f_{N-1,j-1} + b_j f_{N-1,j} + c_j f_{N-1,j+1} = f_{N,j} \quad (20.31)$$

for $j = 1, 2, \dots, M - 1$. The right-hand sides of these equations are known from equation (20.28). Furthermore, from equations (20.29) and (20.30),

$$f_{N-1,0} = K \quad (20.32)$$

$$f_{N-1,M} = 0 \quad (20.33)$$

Equations (20.31) are therefore $M - 1$ simultaneous equations that can be solved for the $M - 1$ unknowns: $f_{N-1,1}, f_{N-1,2}, \dots, f_{N-1,M-1}$.²³ After this has been done, each value

²³ This does not involve inverting a matrix. The $j = 1$ equation in (20.31) can be used to express $f_{N-1,2}$ in terms of $f_{N-1,1}$; the $j = 2$ equation, when combined with the $j = 1$ equation, can be used to express $f_{N-1,3}$ in terms of $f_{N-1,1}$; and so on. The $j = M - 2$ equation, together with earlier equations, enables $f_{N-1,M-1}$ to be expressed in terms of $f_{N-1,1}$. The final $j = M - 1$ equation can then be solved for $f_{N-1,1}$, which can then be used to determine the other $f_{N-1,j}$.

of $f_{N-1,j}$ is compared with $K - j\Delta S$. If $f_{N-1,j} < K - j\Delta S$, early exercise at time $T - \Delta t$ is optimal and $f_{N-1,j}$ is set equal to $K - j\Delta S$. The nodes corresponding to time $T - 2\Delta t$ are handled in a similar way, and so on. Eventually, $f_{0,1}, f_{0,2}, f_{0,3}, \dots, f_{0,M-1}$ are obtained. One of these is the option price of interest.

The control variate technique can be used in conjunction with finite difference methods. The same grid is used to value an option similar to the one under consideration but for which an analytic valuation is available. Equation (20.20) is then used.

Example 20.10

Table 20.4 shows the result of using the implicit finite difference method as just described for pricing the American put option in Example 20.1. Values of 20, 10, and 5 were chosen for M , N , and ΔS , respectively. Thus, the option price is evaluated at \$5 stock price intervals between \$0 and \$100 and at half-month time intervals throughout the life of the option. The option price given by the grid is \$4.07. The same grid gives the price of the corresponding European option as \$3.91. The true European price given by the Black–Scholes–Merton formula is \$4.08. The control variate estimate of the American price is therefore

$$4.07 + (4.08 - 3.91) = \$4.24$$

Explicit Finite Difference Method

The implicit finite difference method has the advantage of being very robust. It always converges to the solution of the differential equation as ΔS and Δt approach zero.²⁴ One of the disadvantages of the implicit finite difference method is that $M - 1$ simultaneous equations have to be solved in order to calculate the $f_{i,j}$ from the $f_{i+1,j}$. The method can be simplified if the values of $\partial f / \partial S$ and $\partial^2 f / \partial S^2$ at point (i, j) on the grid are assumed to be the same as at point $(i + 1, j)$. Equations (20.24) and (20.26) then become

$$\frac{\partial f}{\partial S} = \frac{f_{i+1,j+1} - f_{i+1,j-1}}{2\Delta S}$$

$$\frac{\partial^2 f}{\partial S^2} = \frac{f_{i+1,j+1} + f_{i+1,j-1} - 2f_{i+1,j}}{\Delta S^2}$$

The difference equation is

$$\frac{f_{i+1,j} - f_{i,j}}{\Delta t} + (r - q)j\Delta S \frac{f_{i+1,j+1} - f_{i+1,j-1}}{2\Delta S} + \frac{1}{2}\sigma^2 j^2 \Delta S^2 \frac{f_{i+1,j+1} + f_{i+1,j-1} - 2f_{i+1,j}}{\Delta S^2} = rf_{i,j}$$

or

$$f_{i,j} = a_j^* f_{i+1,j-1} + b_j^* f_{i+1,j} + c_j^* f_{i+1,j+1} \quad (20.34)$$

²⁴ A general rule in finite difference methods is that ΔS should be kept proportional to $\sqrt{\Delta t}$ as they approach zero.

Table 20.4 Grid to value American option in Example 20.1 using implicit finite difference methods.

Stock price (dollars)	Time to maturity (months)										
	5	4.5	4	3.5	3	2.5	2	1.5	1	0.5	0
100	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
95	0.02	0.02	0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00
90	0.05	0.04	0.03	0.02	0.01	0.01	0.00	0.00	0.00	0.00	0.00
85	0.09	0.07	0.05	0.03	0.02	0.01	0.01	0.00	0.00	0.00	0.00
80	0.16	0.12	0.09	0.07	0.04	0.03	0.02	0.01	0.00	0.00	0.00
75	0.27	0.22	0.17	0.13	0.09	0.06	0.03	0.02	0.01	0.00	0.00
70	0.47	0.39	0.32	0.25	0.18	0.13	0.08	0.04	0.02	0.00	0.00
65	0.82	0.71	0.60	0.49	0.38	0.28	0.19	0.11	0.05	0.02	0.00
60	1.42	1.27	1.11	0.95	0.78	0.62	0.45	0.30	0.16	0.05	0.00
55	2.43	2.24	2.05	1.83	1.61	1.36	1.09	0.81	0.51	0.22	0.00
50	4.07	3.88	3.67	3.45	3.19	2.91	2.57	2.17	1.66	0.99	0.00
45	6.58	6.44	6.29	6.13	5.96	5.77	5.57	5.36	5.17	5.02	5.00
40	10.15	10.10	10.05	10.01	10.00	10.00	10.00	10.00	10.00	10.00	10.00
35	15.00	15.00	15.00	15.00	15.00	15.00	15.00	15.00	15.00	15.00	15.00
30	20.00	20.00	20.00	20.00	20.00	20.00	20.00	20.00	20.00	20.00	20.00
25	25.00	25.00	25.00	25.00	25.00	25.00	25.00	25.00	25.00	25.00	25.00
20	30.00	30.00	30.00	30.00	30.00	30.00	30.00	30.00	30.00	30.00	30.00
15	35.00	35.00	35.00	35.00	35.00	35.00	35.00	35.00	35.00	35.00	35.00
10	40.00	40.00	40.00	40.00	40.00	40.00	40.00	40.00	40.00	40.00	40.00
5	45.00	45.00	45.00	45.00	45.00	45.00	45.00	45.00	45.00	45.00	45.00
0	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00

where

$$a_j^* = \frac{1}{1 + r \Delta t} \left(-\frac{1}{2}(r - q)j \Delta t + \frac{1}{2}\sigma^2 j^2 \Delta t \right)$$

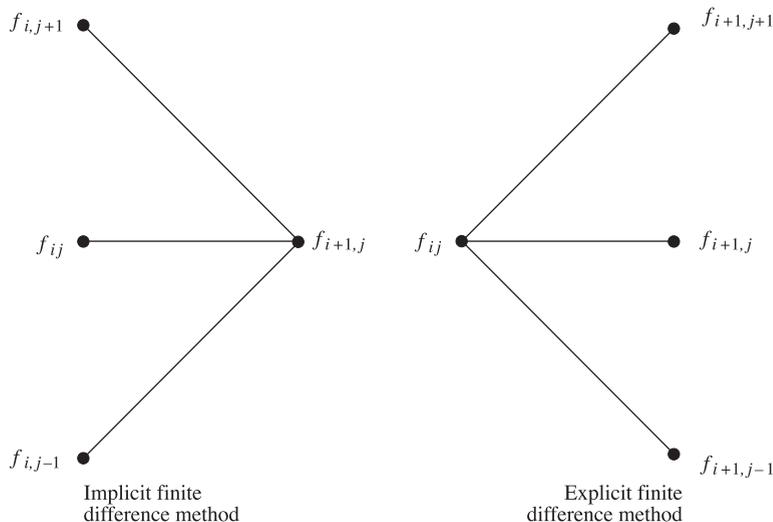
$$b_j^* = \frac{1}{1 + r \Delta t} (1 - \sigma^2 j^2 \Delta t)$$

$$c_j^* = \frac{1}{1 + r \Delta t} \left(\frac{1}{2}(r - q)j \Delta t + \frac{1}{2}\sigma^2 j^2 \Delta t \right)$$

This creates what is known as the *explicit finite difference method*.²⁵ Figure 20.16 shows the difference between the implicit and explicit methods. The implicit method leads to equation (20.27), which gives a relationship between three different values of the option at time $i \Delta t$ (i.e., $f_{i,j-1}$, $f_{i,j}$, and $f_{i,j+1}$) and one value of the option at time $(i + 1) \Delta t$ (i.e., $f_{i+1,j}$). The explicit method leads to equation (20.34), which gives a relationship

²⁵ We also obtain the explicit finite difference method if we use the backward difference approximation instead of the forward difference approximation for $\partial f / \partial t$.

Figure 20.16 Difference between implicit and explicit finite difference methods.



between one value of the option at time $i \Delta t$ (i.e., $f_{i,j}$) and three different values of the option at time $(i + 1) \Delta t$ (i.e., $f_{i+1,j-1}$, $f_{i+1,j}$, $f_{i+1,j+1}$).

Example 20.11

Table 20.5 shows the result of using the explicit version of the finite difference method for pricing the American put option described in Example 20.1. As in Example 20.10, values of 20, 10, and 5 were chosen for M , N , and ΔS , respectively. The option price given by the grid is \$4.26.²⁶

Change of Variable

When geometric Brownian motion is used for the underlying asset price, it is computationally more efficient to use finite difference methods with $\ln S$ rather than S as the underlying variable. Define $Z = \ln S$. Equation (20.21) becomes

$$\frac{\partial f}{\partial t} + \left(r - q - \frac{\sigma^2}{2} \right) \frac{\partial f}{\partial Z} + \frac{1}{2} \sigma^2 \frac{\partial^2 f}{\partial Z^2} = r f$$

The grid then evaluates the derivative for equally spaced values of Z rather than for equally spaced values of S . The difference equation for the implicit method becomes

$$\frac{f_{i+1,j} - f_{i,j}}{\Delta t} + (r - q - \sigma^2/2) \frac{f_{i,j+1} - f_{i,j-1}}{2\Delta Z} + \frac{1}{2} \sigma^2 \frac{f_{i,j+1} + f_{i,j-1} - 2f_{i,j}}{\Delta Z^2} = r f_{i,j}$$

or

$$\alpha_j f_{i,j-1} + \beta_j f_{i,j} + \gamma_j f_{i,j+1} = f_{i+1,j} \tag{20.35}$$

²⁶ The negative numbers and other inconsistencies in the top left-hand part of the grid will be explained later.

Table 20.5 Grid to value American option in Example 20.1 using explicit finite difference methods.

Stock price (dollars)	Time to maturity (months)										
	5	4.5	4	3.5	3	2.5	2	1.5	1	0.5	0
100	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
95	0.06	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
90	-0.11	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
85	0.28	-0.05	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
80	-0.13	0.20	0.00	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00
75	0.46	0.06	0.20	0.04	0.06	0.00	0.00	0.00	0.00	0.00	0.00
70	0.32	0.46	0.23	0.25	0.10	0.09	0.00	0.00	0.00	0.00	0.00
65	0.91	0.68	0.63	0.44	0.37	0.21	0.14	0.00	0.00	0.00	0.00
60	1.48	1.37	1.17	1.02	0.81	0.65	0.42	0.27	0.00	0.00	0.00
55	2.59	2.39	2.21	1.99	1.77	1.50	1.24	0.90	0.59	0.00	0.00
50	4.26	4.08	3.89	3.68	3.44	3.18	2.87	2.53	2.07	1.56	0.00
45	6.76	6.61	6.47	6.31	6.15	5.96	5.75	5.50	5.24	5.00	5.00
40	10.28	10.20	10.13	10.06	10.01	10.00	10.00	10.00	10.00	10.00	10.00
35	15.00	15.00	15.00	15.00	15.00	15.00	15.00	15.00	15.00	15.00	15.00
30	20.00	20.00	20.00	20.00	20.00	20.00	20.00	20.00	20.00	20.00	20.00
25	25.00	25.00	25.00	25.00	25.00	25.00	25.00	25.00	25.00	25.00	25.00
20	30.00	30.00	30.00	30.00	30.00	30.00	30.00	30.00	30.00	30.00	30.00
15	35.00	35.00	35.00	35.00	35.00	35.00	35.00	35.00	35.00	35.00	35.00
10	40.00	40.00	40.00	40.00	40.00	40.00	40.00	40.00	40.00	40.00	40.00
5	45.00	45.00	45.00	45.00	45.00	45.00	45.00	45.00	45.00	45.00	45.00
0	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00

where

$$\alpha_j = \frac{\Delta t}{2\Delta Z}(r - q - \sigma^2/2) - \frac{\Delta t}{2\Delta Z^2}\sigma^2$$

$$\beta_j = 1 + \frac{\Delta t}{\Delta Z^2}\sigma^2 + r\Delta t$$

$$\gamma_j = -\frac{\Delta t}{2\Delta Z}(r - q - \sigma^2/2) - \frac{\Delta t}{2\Delta Z^2}\sigma^2$$

The difference equation for the explicit method becomes

$$\frac{f_{i+1,j} - f_{i,j}}{\Delta t} + (r - q - \sigma^2/2) \frac{f_{i+1,j+1} - f_{i+1,j-1}}{2\Delta Z} + \frac{1}{2}\sigma^2 \frac{f_{i+1,j+1} + f_{i+1,j-1} - 2f_{i+1,j}}{\Delta Z^2} = r f_{i,j}$$

or

$$\alpha_j^* f_{i+1,j-1} + \beta_j^* f_{i+1,j} + \gamma_j^* f_{i+1,j+1} = f_{i,j} \quad (20.36)$$

where

$$\alpha_j^* = \frac{1}{1+r\Delta t} \left[-\frac{\Delta t}{2\Delta Z} (r-q-\sigma^2/2) + \frac{\Delta t}{2\Delta Z^2} \sigma^2 \right] \quad (20.37)$$

$$\beta_j^* = \frac{1}{1+r\Delta t} \left(1 - \frac{\Delta t}{\Delta Z^2} \sigma^2 \right) \quad (20.38)$$

$$\gamma_j^* = \frac{1}{1+r\Delta t} \left[\frac{\Delta t}{2\Delta Z} (r-q-\sigma^2/2) + \frac{\Delta t}{2\Delta Z^2} \sigma^2 \right] \quad (20.39)$$

The change of variable approach has the property that α_j , β_j , and γ_j as well as α_j^* , β_j^* , and γ_j^* are independent of j . In most cases, a good choice for ΔZ is $\sigma\sqrt{3\Delta t}$.

Relation to Trinomial Tree Approaches

The explicit finite difference method is equivalent to the trinomial tree approach.²⁷ In the expressions for a_j^* , b_j^* , and c_j^* in equation (20.34), we can interpret terms as follows:

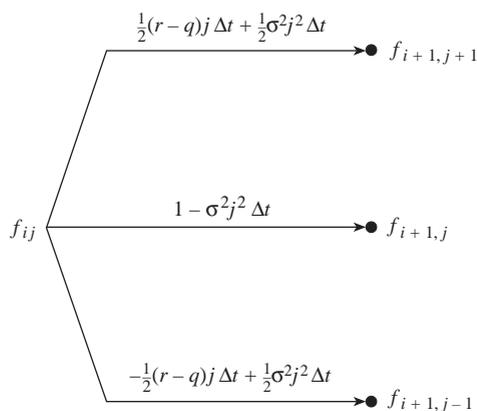
$-\frac{1}{2}(r-q)j\Delta t + \frac{1}{2}\sigma^2 j^2 \Delta t$: Probability of stock price decreasing from $j\Delta S$ to $(j-1)\Delta S$ in time Δt .

$1 - \sigma^2 j^2 \Delta t$: Probability of stock price remaining unchanged at $j\Delta S$ in time Δt .

$\frac{1}{2}(r-q)j\Delta t + \frac{1}{2}\sigma^2 j^2 \Delta t$: Probability of stock price increasing from $j\Delta S$ to $(j+1)\Delta S$ in time Δt .

This interpretation is illustrated in Figure 20.17. The three probabilities sum to unity. They give the expected increase in the stock price in time Δt as $(r-q)j\Delta S\Delta t = (r-q)S\Delta t$. This is the expected increase in a risk-neutral world. For small values

Figure 20.17 Interpretation of explicit finite difference method as a trinomial tree.



²⁷ It can also be shown that the implicit finite difference method is equivalent to a multinomial tree approach where there are $M+1$ branches emanating from each node.

of Δt , they also give the variance of the change in the stock price in time Δt as $\sigma^2 j^2 \Delta S^2 \Delta t = \sigma^2 S^2 \Delta t$. This corresponds to the stochastic process followed by S . The value of f at time $i \Delta t$ is calculated as the expected value of f at time $(i + 1) \Delta t$ in a risk-neutral world discounted at the risk-free rate.

For the explicit version of the finite difference method to work well, the three “probabilities”

$$\begin{aligned} &-\frac{1}{2}(r - q)j \Delta t + \frac{1}{2}\sigma^2 j^2 \Delta t, \\ &1 - \sigma^2 j^2 \Delta t \\ &\frac{1}{2}(r - q)j \Delta t + \frac{1}{2}\sigma^2 j^2 \Delta t \end{aligned}$$

should all be positive. In Example 20.11, $1 - \sigma^2 j^2 \Delta t$ is negative when $j \geq 13$ (i.e., when $S \geq 65$). This explains the negative option prices and other inconsistencies in the top left-hand part of Table 20.5. This example illustrates the main problem associated with the explicit finite difference method. Because the probabilities in the associated tree may be negative, it does not necessarily produce results that converge to the solution of the differential equation.²⁸

When the change-of-variable approach is used (see equations (20.36) to (20.39)), the probability that $Z = \ln S$ will decrease by ΔZ , stay the same, and increase by ΔZ are

$$\begin{aligned} &-\frac{\Delta t}{2\Delta Z}(r - q - \sigma^2/2) + \frac{\Delta t}{2\Delta Z^2}\sigma^2 \\ &1 - \frac{\Delta t}{\Delta Z^2}\sigma^2 \\ &\frac{\Delta t}{2\Delta Z}(r - q - \sigma^2/2) + \frac{\Delta t}{2\Delta Z^2}\sigma^2 \end{aligned}$$

respectively. These movements in Z correspond to the stock price changing from S to $Se^{-\Delta Z}$, S , and $Se^{\Delta Z}$, respectively. If we set $\Delta Z = \sigma\sqrt{3\Delta t}$, then the tree and the probabilities are identical to those for the trinomial tree approach discussed in Section 20.4.

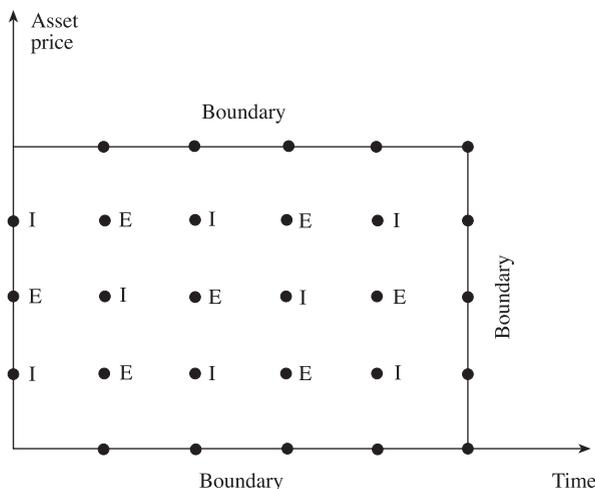
Other Finite Difference Methods

Many of the other finite difference methods that have been proposed have some of the features of the explicit finite difference method and some features of the implicit finite difference method.

In what is known as the *hopscotch method*, we alternate between the explicit and implicit calculations as we move from node to node. This is illustrated in Figure 20.18. At each time, we first do all the calculations at the “explicit nodes” (E) in the usual way. The “implicit nodes” (I) can then be handled without solving a set of simultaneous equations because the values at the adjacent nodes have already been calculated.

²⁸ J. Hull and A. White, “Valuing Derivative Securities Using the Explicit Finite Difference Method,” *Journal of Financial and Quantitative Analysis*, 25 (March 1990): 87–100, show how this problem can be overcome. In the situation considered here it is sufficient to construct the grid in $\ln S$ rather than S to ensure convergence.

Figure 20.18 The hopscotch method. I indicates node at which implicit calculations are done; E indicates node at which explicit calculations are done.



The Crank–Nicolson method is an average of the explicit and implicit methods. For the implicit method, equation (20.27) gives

$$f_{i,j} = a_j f_{i-1,j-1} + b_j f_{i-1,j} + c_j f_{i-1,j+1}$$

For the explicit method, equation (20.34) gives

$$f_{i-1,j} = a_j^* f_{i,j-1} + b_j^* f_{i,j} + c_j^* f_{i,j+1}$$

The Crank–Nicolson method averages these two equations to obtain

$$f_{i,j} + f_{i-1,j} = a_j f_{i-1,j-1} + b_j f_{i-1,j} + c_j f_{i-1,j+1} + a_j^* f_{i,j-1} + b_j^* f_{i,j} + c_j^* f_{i,j+1}$$

Putting

$$g_{i,j} = f_{i,j} - a_j^* f_{i,j-1} - b_j^* f_{i,j} - c_j^* f_{i,j+1}$$

gives

$$g_{i,j} = a_j f_{i-1,j-1} + b_j f_{i-1,j} + c_j f_{i-1,j+1} - f_{i-1,j}$$

This shows that implementing the Crank–Nicolson method is similar to implementing the implicit finite difference method. The advantage of the Crank–Nicolson method is that it has faster convergence than either the explicit or implicit method.

Applications of Finite Difference Methods

Finite difference methods can be used for the same types of derivative pricing problems as tree approaches. They can handle American-style as well as European-style derivatives but cannot easily be used in situations where the payoff from a derivative depends on the past history of the underlying variable. Finite difference methods can, at the expense of a considerable increase in computer time, be used when there are several state variables. The grid in Figure 20.15 then becomes multidimensional.

The method for calculating Greek letters is similar to that used for trees. Delta, gamma, and theta can be calculated directly from the $f_{i,j}$ values on the grid. For vega, it is necessary to make a small change to volatility and recalculate the value of the derivative using the same grid.

SUMMARY

We have presented three different numerical procedures for valuing derivatives when no analytic solution is available. These involve the use of trees, Monte Carlo simulation, and finite difference methods.

Binomial trees assume that, in each short interval of time Δt , a stock price either moves up by a multiplicative amount u or down by a multiplicative amount d . The sizes of u and d and their associated probabilities are chosen so that the change in the stock price has the correct mean and standard deviation in a risk-neutral world. Derivative prices are calculated by starting at the end of the tree and working backwards. For an American option, the value at a node is the greater of (a) the value if it is exercised immediately and (b) the discounted expected value if it is held for a further period of time Δt .

Monte Carlo simulation involves using random numbers to sample many different paths that the variables underlying the derivative could follow in a risk-neutral world. For each path, the payoff is calculated and discounted at the risk-free interest rate. The arithmetic average of the discounted payoffs is the estimated value of the derivative.

Finite difference methods solve the underlying differential equation by converting it to a difference equation. They are similar to tree approaches in that the computations work back from the end of the life of the derivative to the beginning. The explicit finite difference method is functionally the same as using a trinomial tree. The implicit finite difference method is more complicated but has the advantage that the user does not have to take any special precautions to ensure convergence.

In practice, the method that is chosen is likely to depend on the characteristics of the derivative being evaluated and the accuracy required. Monte Carlo simulation works forward from the beginning to the end of the life of a derivative. It can be used for European-style derivatives and can cope with a great deal of complexity as far as the payoffs are concerned. It becomes relatively more efficient as the number of underlying variables increases. Tree approaches and finite difference methods work from the end of the life of a security to the beginning and can accommodate American-style as well as European-style derivatives. However, they are difficult to apply when the payoffs depend on the past history of the state variables as well as on their current values. Also, they are liable to become computationally very time consuming when three or more variables are involved.

FURTHER READING

General

Clewlow, L., and C. Strickland, *Implementing Derivatives Models*. Chichester: Wiley, 1998.

Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*, 2nd edn. Cambridge University Press, 1992.

On Tree Approaches

Cox, J. C., S. A. Ross, and M. Rubinstein. "Option Pricing: A Simplified Approach," *Journal of Financial Economics*, 7 (October 1979): 229–64.

Figlewski, S., and B. Gao. "The Adaptive Mesh Model: A New Approach to Efficient Option Pricing," *Journal of Financial Economics*, 53 (1999): 313–51.

Hull, J. C., and A. White, "The Use of the Control Variate Technique in Option Pricing," *Journal of Financial and Quantitative Analysis*, 23 (September 1988): 237–51.

Rendleman, R., and B. Bartter, "Two State Option Pricing," *Journal of Finance*, 34 (1979): 1092–1110.

On Monte Carlo Simulation

Boyle, P. P., "Options: A Monte Carlo Approach," *Journal of Financial Economics*, 4 (1977): 323–38.

Boyle, P. P., M. Broadie, and P. Glasserman. "Monte Carlo Methods for Security Pricing," *Journal of Economic Dynamics and Control*, 21 (1997): 1267–1322.

Broadie, M., P. Glasserman, and G. Jain. "Enhanced Monte Carlo Estimates for American Option Prices," *Journal of Derivatives*, 5 (Fall 1997): 25–44.

On Finite Difference Methods

Hull, J. C., and A. White, "Valuing Derivative Securities Using the Explicit Finite Difference Method," *Journal of Financial and Quantitative Analysis*, 25 (March 1990): 87–100.

Wilmott, P., *Derivatives: The Theory and Practice of Financial Engineering*. Chichester: Wiley, 1998.

Practice Questions (Answers in Solutions Manual)

- 20.1. Which of the following can be estimated for an American option by constructing a single binomial tree: delta, gamma, vega, theta, rho?
- 20.2. Calculate the price of a 3-month American put option on a non-dividend-paying stock when the stock price is \$60, the strike price is \$60, the risk-free interest rate is 10% per annum, and the volatility is 45% per annum. Use a binomial tree with a time interval of 1 month.
- 20.3. Explain how the control variate technique is implemented when a tree is used to value American options.
- 20.4. Calculate the price of a 9-month American call option on corn futures when the current futures price is 198 cents, the strike price is 200 cents, the risk-free interest rate is 8% per annum, and the volatility is 30% per annum. Use a binomial tree with a time interval of 3 months.
- 20.5. Consider an option that pays off the amount by which the final stock price exceeds the average stock price achieved during the life of the option. Can this be valued using the binomial tree approach? Explain your answer.
- 20.6. "For a dividend-paying stock, the tree for the stock price does not recombine; but the tree for the stock price less the present value of future dividends does recombine." Explain this statement.
- 20.7. Show that the probabilities in a Cox, Ross, and Rubinstein binomial tree are negative when the condition in footnote 9 holds.
- 20.8. Use stratified sampling with 100 trials to improve the estimate of π in Business Snapshot 20.1 and Table 20.1.

- 20.9. Explain why the Monte Carlo simulation approach cannot easily be used for American-style derivatives.
- 20.10. A 9-month American put option on a non-dividend-paying stock has a strike price of \$49. The stock price is \$50, the risk-free rate is 5% per annum, and the volatility is 30% per annum. Use a three-step binomial tree to calculate the option price.
- 20.11. Use a three-time-step tree to value a 9-month American call option on wheat futures. The current futures price is 400 cents, the strike price is 420 cents, the risk-free rate is 6%, and the volatility is 35% per annum. Estimate the delta of the option from your tree.
- 20.12. A 3-month American call option on a stock has a strike price of \$20. The stock price is \$20, the risk-free rate is 3% per annum, and the volatility is 25% per annum. A dividend of \$2 is expected in 1.5 months. Use a three-step binomial tree to calculate the option price.
- 20.13. A 1-year American put option on a non-dividend-paying stock has an exercise price of \$18. The current stock price is \$20, the risk-free interest rate is 15% per annum, and the volatility of the stock price is 40% per annum. Use the DerivaGem software with four 3-month time steps to estimate the value of the option. Display the tree and verify that the option prices at the final and penultimate nodes are correct. Use DerivaGem to value the European version of the option. Use the control variate technique to improve your estimate of the price of the American option.
- 20.14. A 2-month American put option on a stock index has an exercise price of 480. The current level of the index is 484, the risk-free interest rate is 10% per annum, the dividend yield on the index is 3% per annum, and the volatility of the index is 25% per annum. Divide the life of the option into four half-month periods and use the tree approach to estimate the value of the option.
- 20.15. How can the control variate approach improve the estimate of the delta of an American option when the tree approach is used?
- 20.16. Suppose that Monte Carlo simulation is being used to evaluate a European call option on a non-dividend-paying stock when the volatility is stochastic. How could the control variate and antithetic variable technique be used to improve numerical efficiency? Explain why it is necessary to calculate six values of the option in each simulation trial when both the control variate and the antithetic variable technique are used.
- 20.17. Explain how equations (20.27) to (20.30) change when the implicit finite difference method is being used to evaluate an American call option on a currency.
- 20.18. An American put option on a non-dividend-paying stock has 4 months to maturity. The exercise price is \$21, the stock price is \$20, the risk-free rate of interest is 10% per annum, and the volatility is 30% per annum. Use the explicit version of the finite difference approach to value the option. Use stock price intervals of \$4 and time intervals of 1 month.
- 20.19. The spot price of copper is \$0.60 per pound. Suppose that the futures prices (dollars per pound) are as follows:

3 months	0.59
6 months	0.57
9 months	0.54
12 months	0.50

The volatility of the price of copper is 40% per annum and the risk-free rate is 6% per

- annum. Use a binomial tree to value an American call option on copper with an exercise price of \$0.60 and a time to maturity of 1 year. Divide the life of the option into four 3-month periods for the purposes of constructing the tree. (*Hint*: As explained in Section 17.7, the futures price of a variable is its expected future price in a risk-neutral world.)
- 20.20. Use the binomial tree in Problem 20.19 to value a security that pays off x^2 in 1 year where x is the price of copper.
- 20.21. When do the boundary conditions for $S = 0$ and $S \rightarrow \infty$ affect the estimates of derivative prices in the explicit finite difference method?
- 20.22. How would you use the antithetic variable method to improve the estimate of the European option in Business Snapshot 20.2 and Table 20.2?
- 20.23. A company has issued a 3-year convertible bond that has a face value of \$25 and can be exchanged for two of the company's shares at any time. The company can call the issue, forcing conversion, when the share price is greater than or equal to \$18. Assuming that the company will force conversion at the earliest opportunity, what are the boundary conditions for the price of the convertible? Describe how you would use finite difference methods to value the convertible assuming constant interest rates. Assume there is no risk of the company defaulting.
- 20.24. Provide formulas that can be used for obtaining three random samples from standard normal distributions when the correlation between sample i and sample j is $\rho_{i,j}$.

Further Questions

- 20.25. An American put option to sell a Swiss franc for dollars has a strike price of \$0.80 and a time to maturity of 1 year. The Swiss franc's volatility is 10%, the dollar interest rate is 6%, the Swiss franc interest rate is 3%, and the current exchange rate is 0.81. Use a three-step binomial tree to value the option. Estimate the delta of the option from your tree.
- 20.26. A 1-year American call option on silver futures has an exercise price of \$9.00. The current futures price is \$8.50, the risk-free rate of interest is 12% per annum, and the volatility of the futures price is 25% per annum. Use the DerivaGem software with four 3-month time steps to estimate the value of the option. Display the tree and verify that the option prices at the final and penultimate nodes are correct. Use DerivaGem to value the European version of the option. Use the control variate technique to improve your estimate of the price of the American option.
- 20.27. A 6-month American call option on a stock is expected to pay dividends of \$1 per share at the end of the second month and the fifth month. The current stock price is \$30, the exercise price is \$34, the risk-free interest rate is 10% per annum, and the volatility of the part of the stock price that will not be used to pay the dividends is 30% per annum. Use the DerivaGem software with the life of the option divided into six time steps to estimate the value of the option. Compare your answer with that given by Black's approximation (see Section 14.12).
- 20.28. The current value of the British pound is \$1.60 and the volatility of the pound/dollar exchange rate is 15% per annum. An American call option has an exercise price of \$1.62 and a time to maturity of 1 year. The risk-free rates of interest in the United States and

- the United Kingdom are 6% per annum and 9% per annum, respectively. Use the explicit finite difference method to value the option. Consider exchange rates at intervals of 0.20 between 0.80 and 2.40 and time intervals of 3 months.
- 20.29. Answer the following questions concerned with the alternative procedures for constructing trees in Section 20.4:
- Show that the binomial model in Section 20.4 is exactly consistent with the mean and variance of the change in the logarithm of the stock price in time Δt .
 - Show that the trinomial model in Section 20.4 is consistent with the mean and variance of the change in the logarithm of the stock price in time Δt when terms of order $(\Delta t)^2$ and higher are ignored.
 - Construct an alternative to the trinomial model in Section 20.4 so that the probabilities are $1/6$, $2/3$, and $1/6$ on the upper, middle, and lower branches emanating from each node. Assume that the branching is from S to Su , Sm , or Sd with $m^2 = ud$. Match the mean and variance of the change in the logarithm of the stock price exactly.
- 20.30. The DerivaGem Application Builder functions enable you to investigate how the prices of options calculated from a binomial tree converge to the correct value as the number of time steps increases. (See Figure 20.4 and Sample Application A in DerivaGem.) Consider a put option on a stock index where the index level is 900, the strike price is 900, the risk-free rate is 5%, the dividend yield is 2%, and the time to maturity is 2 years.
- Produce results similar to Sample Application A on convergence for the situation where the option is European and the volatility of the index is 20%.
 - Produce results similar to Sample Application A on convergence for the situation where the option is American and the volatility of the index is 20%.
 - Produce a chart showing the pricing of the American option when the volatility is 20% as a function of the number of time steps when the control variate technique is used.
 - Suppose that the price of the American option in the market is 85.0. Produce a chart showing the implied volatility estimate as a function of the number of time steps.
- 20.31. Estimate delta, gamma, and theta from the tree in Example 20.3. Explain how each can be interpreted.
- 20.32. How much is gained from exercising early at the lowest node at the 9-month point in Example 20.4?

Business Snapshot 21.1 How Bank Regulators Use VaR

The Basel Committee on Bank Supervision is a committee of the world's bank regulators that meets regularly in Basel, Switzerland. In 1988 it published what has become known as Basel I. This is an agreement between the regulators on how the capital a bank is required to hold for credit risk should be calculated. Later the Basel Committee published *The 1996 Amendment*, which was implemented in 1998 and required banks to hold capital for market risk as well as credit risk. The amendment distinguished between a bank's trading book and its banking book. The banking book consists primarily of loans and is not usually revalued on a regular basis for managerial and accounting purposes. The trading book consists of the myriad of different instruments that are traded by the bank (stocks, bonds, swaps, forward contracts, options, etc.) and is normally revalued daily.

The 1996 amendment calculated capital for the trading book using the VaR measure with $N = 10$ and $X = 99$. This means that it focused on the revaluation loss over a 10-day period that is expected to be exceeded only 1% of the time. The capital it required the bank to hold is k times this VaR measure (with an adjustment for what are termed specific risks). The multiplier k was chosen on a bank-by-bank basis by the regulators and must be at least 3.0. For a bank with excellent well-tested VaR estimation procedures, it was likely that k would be set equal to the minimum value of 3.0. For other banks, it could be higher. Following the credit crisis that started in 2007, the rules were revised.

a negative loss and VaR is concerned with the right tail of the distribution.) For example, when $N = 5$ and $X = 97$, VaR is the third percentile of the distribution of gain in the value of the portfolio over the next 5 days. VaR is illustrated for the situation where the change in the value of the portfolio is approximately normally distributed in Figure 21.1.

VaR is an attractive measure because it is easy to understand. In essence, it asks the simple question "How bad can things get?" This is the question all senior managers want answered. They are very comfortable with the idea of compressing all the Greek letters for all the market variables underlying a portfolio into a single number.

If we accept that it is useful to have a single number to describe the risk of a portfolio, an interesting question is whether VaR is the best alternative. Some researchers have argued that VaR may tempt traders to choose a portfolio with a return distribution similar to that in Figure 21.2. The portfolios in Figures 21.1 and 21.2 have the same VaR, but the portfolio in Figure 21.2 is much riskier because potential losses are much larger.

A measure that deals with the problem we have just mentioned is *expected shortfall*.¹ Whereas VaR asks the question "How bad can things get?", expected shortfall asks "If things do get bad, how much can the company expect to lose?" Expected shortfall is the expected loss during an N -day period conditional that an outcome in

¹ This measure, which is also known as *C-VaR* or *tail loss*, was suggested by P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath, "Coherent Measures of Risk," *Mathematical Finance*, 9 (1999): 203–28. These authors define certain properties that a good risk measure should have and show that the standard VaR measure does not have all of them. For more details, see J. Hull, *Risk Management and Financial Institutions*, 2nd edn., 2010.

Figure 21.1 Calculation of VaR from the probability distribution of the change in the portfolio value; confidence level is $X\%$. Gains in portfolio value are positive; losses are negative.

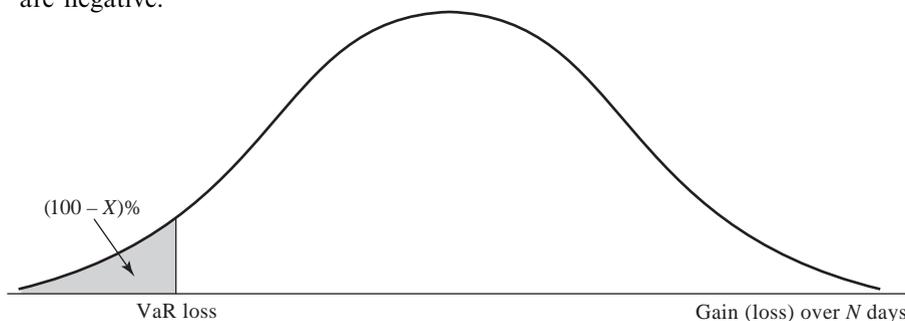
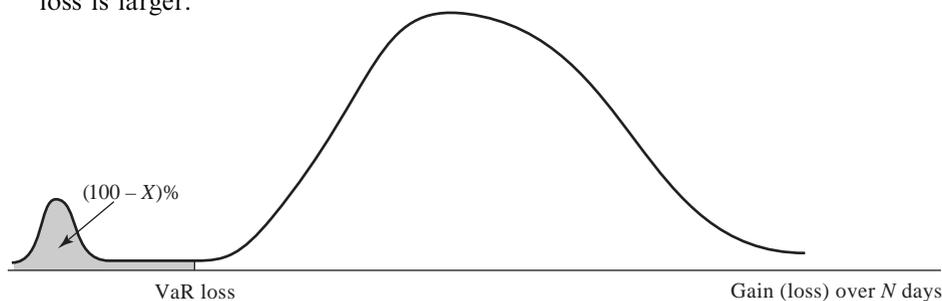


Figure 21.2 Alternative situation to Figure 21.1. VaR is the same, but the potential loss is larger.



the $(100 - X)\%$ left tail of the distribution occurs. For example, with $X = 99$ and $N = 10$, the expected shortfall is the average amount the company loses over a 10-day period when the loss is in the 1% tail of the distribution.

In spite of its weaknesses, VaR (not expected shortfall) is the most popular measure of risk among both regulators and risk managers. We will therefore devote most of the rest of this chapter to how it can be measured.

The Time Horizon

VaR has two parameters: the time horizon N , measured in days, and the confidence level X . In practice, analysts almost invariably set $N = 1$ in the first instance. This is because there is not enough data to estimate directly the behavior of market variables over periods of time longer than 1 day. The usual assumption is

$$N\text{-day VaR} = 1\text{-day VaR} \times \sqrt{N}$$

This formula is exactly true when the changes in the value of the portfolio on successive days have independent identical normal distributions with mean zero. In other cases it is an approximation.

Business Snapshot 21.1 explains that the 1996 amendment to Basel I required a bank's capital for market risk to be at least three times the 10-day 99% VaR. Given the way a 10-day VaR is calculated, this is $3 \times \sqrt{10} = 9.49$ times the 1-day 99% VaR.

21.2 HISTORICAL SIMULATION

Historical simulation is one popular way of estimating VaR. It involves using past data as a guide to what will happen in the future. Suppose that we want to calculate VaR for a portfolio using a one-day time horizon, a 99% confidence level, and 501 days of data. (The time horizon and confidence level are those typically used for a market risk VaR calculation; 501 is a popular choice for the number of days of data used because, as we shall see, it leads to 500 scenarios being created.) The first step is to identify the market variables affecting the portfolio. These will typically be interest rates, equity prices, commodity prices, and so on. All prices are measured in the domestic currency. For example, one market variable for a German bank is likely to be the S&P 500 measured in euros.

Data are collected on movements in the market variables over the most recent 501 days. This provides 500 alternative scenarios for what can happen between today and tomorrow. Denote the first day for which we have data as Day 0, the second day as Day 1, and so on. Scenario 1 is where the percentage changes in the values of all variables are the same as they were between Day 0 and Day 1, Scenario 2 is where they are the same as between Day 1 and Day 2, and so on. For each scenario, the dollar change in the value of the portfolio between today and tomorrow is calculated. This defines a probability distribution for daily loss (gains are negative losses) in the value of our portfolio. The 99th percentile of the distribution can be estimated as the fifth-highest loss.² The estimate of VaR is the loss when we are at this 99th percentile point. We are 99% certain that we will not take a loss greater than the VaR estimate if the changes in market variables in the last 501 days are representative of what will happen between today and tomorrow.

To express the approach algebraically, define v_i as the value of a market variable on Day i and suppose that today is Day n . The i th scenario in the historical simulation approach assumes that the value of the market variable tomorrow will be

$$\text{Value under } i\text{th scenario} = v_n \frac{v_i}{v_{i-1}}$$

Illustration: Investment in Four Stock Indices

To illustrate the calculations underlying the approach, suppose that an investor in the United States owns, on September 25, 2008, a portfolio worth \$10 million consisting of investments in four stock indices: the Dow Jones Industrial Average (DJIA) in the US, the FTSE 100 in the UK, the CAC 40 in France, and the Nikkei 225 in Japan. The value of the investment in each index on September 25, 2008, is shown in Table 21.1. An Excel spreadsheet containing 501 days of historical data on the closing prices of the

² There are alternatives here. A case can be made for using the fifth-highest loss, the sixth-highest loss, or an average of the two. In Excel's PERCENTILE function, when there are n observations and k is an integer, the $k/(n - 1)$ percentile is the observation ranked $k + 1$. Other percentiles are calculated using linear interpolation.

Table 21.1 Investment portfolio used for VaR calculations.

<i>Index</i>	<i>Portfolio value (\$000s)</i>
DJIA	\$4,000
FTSE 100	\$3,000
CAC 40	\$1,000
Nikkei 225	\$2,000
<i>Total</i>	\$10,000

four indices, together with exchange rates and a complete set of VaR calculations are on the author's website:³

www.rotman.utoronto.ca/~hull/OFOD/VaRExample

The quoted values of the FTSE 100, CAC 40, and Nikkei 225 are adjusted for exchange rate changes so that they are measured in US dollars. For example, the FTSE 100 was 5197.00 on September 25, 2008, when the exchange rate was 1.8472 USD per GBP. It was 5823.40 on August 10, 2006, when the exchange rate was 1.8918 USD per GBP. When measuring in USD, if the index is set to 5197.00 on September 25, 2008 it is

$$5,823.40 \times \frac{1.8918}{1.8472} = 5,964.00$$

on August 10, 2006. An extract from the data after exchange rate adjustments have been made is shown in Table 21.2.

September 25, 2008, is an interesting date to choose in evaluating an equity investment. The turmoil in credit markets, which started in August 2007, was over a year old. Equity prices had been declining for several months. Volatilities were increasing. Lehman Brothers had filed for bankruptcy ten days earlier. The Treasury Secretary's \$700 billion Troubled Asset Relief Program (TARP) had not yet been passed by the United States Congress.

Table 21.2 Data on stock indices for historical simulation after exchange rate adjustments.

<i>Day</i>	<i>Date</i>	<i>DJIA</i>	<i>FTSE 100</i>	<i>CAC 40</i>	<i>Nikkei 225</i>
0	Aug. 7, 2006	11,219.38	6,026.33	4,345.08	14,023.44
1	Aug. 8, 2006	11,173.59	6,007.08	4,347.99	14,300.91
2	Aug. 9, 2006	11,076.18	6,055.30	4,413.35	14,467.09
3	Aug. 10, 2006	11,124.37	5,964.00	4,333.90	14,413.32
⋮	⋮	⋮	⋮	⋮	⋮
499	Sept. 24, 2008	10,825.17	5,109.67	4,113.33	12,159.59
500	Sept. 25, 2008	11,022.06	5,197.00	4,226.81	12,006.53

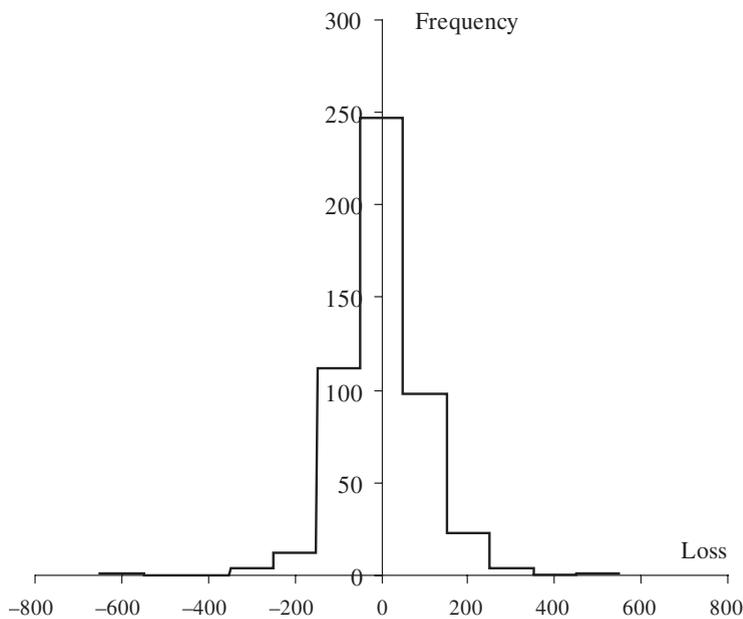
³ To keep the example as straightforward as possible, only days when all four indices traded were included in the compilation of the data. This is why the 501 items of data extend from August 7, 2006 to September 25, 2008. In practice, if the analysis were carried out by a US financial institution, an attempt might well be made to fill in data for days that were not US holidays.

Table 21.3 Scenarios generated for September 26, 2008, using data in Table 21.2.

Scenario number	DJIA	FTSE 100	CAC 40	Nikkei 225	Portfolio value (\$000s)	Loss (\$000s)
1	10,977.08	5,180.40	4,229.64	12,244.10	10,014.334	-14.334
2	10,925.97	5,238.72	4,290.35	12,146.04	10,027.481	-27.481
3	11,070.01	5,118.64	4,150.71	11,961.91	9,946.736	53.264
⋮	⋮	⋮	⋮	⋮	⋮	⋮
499	10,831.43	5,079.84	4,125.61	12,115.90	9,857.465	142.535
500	11,222.53	5,285.82	4,343.42	11,855.40	10,126.439	-126.439

Table 21.3 shows the values of the market variables on September 26, 2008, for the scenarios considered. Scenario 1 (the first row in Table 21.3) shows the values of market variables on September 26, 2008, assuming that their percentage changes between September 25 and September 26, 2008, are the same as they were between August 7 and August 8, 2006; Scenario 2 (the second row in Table 21.3) shows the values of market variables on September 26, 2008, assuming these percentage changes are the same as those between August 8 and August 9, 2006; and so on. In general, Scenario i assumes that the percentage changes in the indices between September 25 and September 26 are the same as they were between Day $i - 1$ and Day i for $1 \leq i \leq 500$. The 500 rows in Table 21.3 are the 500 scenarios considered.

The DJIA was 11,022.06 on September 25, 2008. On August 8, 2006, it was 11,173.59, down from 11,219.38 on August 7, 2006. Therefore the value of the DJIA under

Figure 21.3 Histogram of losses for the scenarios considered between September 25 and September 26, 2008.

Scenario 1 is

$$11,022.06 \times \frac{11,173.59}{11,219.38} = 10,977.08$$

Similarly, the values of the FTSE 100, the CAC 40, and the Nikkei 225 are 5,180.40, 4,229.64, and 12,244.10, respectively. Therefore the value of the portfolio under Scenario 1 is (in \$000s)

$$\begin{aligned} 4,000 \times \frac{10,977.08}{11,022.06} + 3,000 \times \frac{5,180.40}{5,197.00} \\ + 1,000 \times \frac{4,229.64}{4,226.81} + 2,000 \times \frac{12,224.10}{12,006.53} = 10,014.334 \end{aligned}$$

The portfolio therefore has a gain of \$14,334 under Scenario 1. A similar calculation is carried out for the other scenarios. A histogram for the losses is shown in Figure 21.3. (The bars on the histogram represent losses (\$000s) in the ranges 450 to 550, 350 to 450, 250 to 350, and so on.)

The losses for the 500 different scenarios are then ranked. An extract from the results of doing this is shown in Table 21.4. The worst scenario is number 494 (where indices are assumed to change in the same way that they did at the time of the bankruptcy of Lehman Brothers). The one-day 99% value at risk can be estimated as the fifth-worst loss. This is \$253,385.

As explained in Section 21.1, the ten-day 99% VaR is usually calculated as $\sqrt{10}$ times the one-day 99% VaR. In this case the ten-day VaR would therefore be

$$\sqrt{10} \times 253,385 = 801,274$$

or \$801,274.

Table 21.4 Losses ranked from highest to lowest for 500 scenarios.

<i>Scenario number</i>	<i>Loss (\$000s)</i>
494	477.841
339	345.435
349	282.204
329	277.041
487	253.385
227	217.974
131	205.256
238	201.389
473	191.269
306	191.050
477	185.127
495	184.450
376	182.707
237	180.105
365	172.224
⋮	⋮

Each day the VaR estimate in our example would be updated using the most recent 501 days of data. Consider, for example, what happens on September 26, 2008 (Day 501). We find out new values for all the market variables and are able to calculate a new value for our portfolio. We then go through the procedure we have outlined to calculate a new VaR. Data on the market variables from August 8, 2006, to September 26, 2008 (Day 1 to Day 501) are used in the calculation. (This gives us the required 500 observations on the percentage changes in market variables; the August 7, 2006, Day 0, values of the market variables are no longer used.) Similarly, on the next trading day September 29, 2008 (Day 502), data from August 9, 2006, to September 29, 2008 (Day 2 to Day 502) are used to determine VaR, and so on.

In practice, a financial institution's portfolio is, of course, considerably more complicated than the one we have considered here. It is likely to consist of thousands or tens of thousands of positions. Some of the bank's positions are typically in forward contracts, options, and other derivatives. The VaR is calculated at the end of each day on the assumption that the portfolio will remain unchanged over the next business day. If a bank's trading during a day leads to a more risky (less risky) portfolio, the ten-day 99% VaR typically increases (decreases) over the previous day's value.

It is often necessary to consider hundreds or even thousands of market variables in a VaR calculation. In the case of interest rates, a bank typically needs the Treasury and LIBOR/swap term structure of zero-coupon interest rates in a number of different currencies in order to value its portfolio. The market variables that are considered are the ones from which these term structures are calculated (see Chapter 4 for the calculation of the term structure of zero rates). There might be as many as ten market variables for each zero curve to which the bank is exposed.

21.3 MODEL-BUILDING APPROACH

The main alternative to historical simulation is the model-building approach. Before getting into the details of the approach, it is appropriate to mention one issue concerned with the units for measuring volatility.

Daily Volatilities

In option pricing, time is usually measured in years, and the volatility of an asset is usually quoted as a "volatility per year". When using the model-building approach to calculate VaR, time is usually measured in days and the volatility of an asset is usually quoted as a "volatility per day."

What is the relationship between the volatility per year used in option pricing and the volatility per day used in VaR calculations? Let us define σ_{year} as the volatility per year of a certain asset and σ_{day} as the equivalent volatility per day of the asset. Assuming 252 trading days in a year, equation (14.2) gives the standard deviation of the continuously compounded return on the asset in 1 year as either σ_{year} or $\sigma_{\text{day}}\sqrt{252}$. It follows that

$$\sigma_{\text{year}} = \sigma_{\text{day}}\sqrt{252}$$

or

$$\sigma_{\text{day}} = \frac{\sigma_{\text{year}}}{\sqrt{252}}$$

so that daily volatility is about 6% of annual volatility.

As pointed out in Section 14.4, σ_{day} is approximately equal to the standard deviation of the percentage change in the asset price in one day. For the purposes of calculating VaR we assume exact equality. The daily volatility of an asset price (or any other variable) is therefore defined as equal to the standard deviation of the percentage change in one day.

Our discussion in the next few sections assumes that estimates of daily volatilities and correlations are available. Chapter 22 discusses how the estimates can be produced.

Single-Asset Case

Consider how VaR is calculated using the model-building approach in a very simple situation where the portfolio consists of a position in a single stock: \$10 million in shares of Microsoft. We suppose that $N = 10$ and $X = 99$, so that we are interested in the loss level over 10 days that we are 99% confident will not be exceeded. Initially, we consider a 1-day time horizon.

Assume that the volatility of Microsoft is 2% per day (corresponding to about 32% per year). Because the size of the position is \$10 million, the standard deviation of daily changes in the value of the position is 2% of \$10 million, or \$200,000.

It is customary in the model-building approach to assume that the expected change in a market variable over the time period considered is zero. This is not strictly true, but it is a reasonable assumption. The expected change in the price of a market variable over a short time period is generally small when compared with the standard deviation of the change. Suppose, for example, that Microsoft has an expected return of 20% per annum. Over a 1-day period, the expected return is $0.20/252$, or about 0.08%, whereas the standard deviation of the return is 2%. Over a 10-day period, the expected return is 0.08×10 , or about 0.8%, whereas the standard deviation of the return is $2\sqrt{10}$, or about 6.3%.

So far, we have established that the change in the value of the portfolio of Microsoft shares over a 1-day period has a standard deviation of \$200,000 and (at least approximately) a mean of zero. We assume that the change is normally distributed.⁴ From the tables at the end of this book, $N(-2.33) = 0.01$. This means that there is a 1% probability that a normally distributed variable will decrease in value by more than 2.33 standard deviations. Equivalently, it means that we are 99% certain that a normally distributed variable will not decrease in value by more than 2.33 standard deviations. The 1-day 99% VaR for our portfolio consisting of a \$10 million position in Microsoft is therefore

$$2.33 \times 200,000 = \$466,000$$

As discussed earlier, the N -day VaR is calculated as \sqrt{N} times the 1-day VaR. The 10-day 99% VaR for Microsoft is therefore

$$466,000 \times \sqrt{10} = \$1,473,621$$

Consider next a portfolio consisting of a \$5 million position in AT&T, and suppose the daily volatility of AT&T is 1% (approximately 16% per year). A similar calculation

⁴ To be consistent with the option pricing assumption in Chapter 14, we could assume that the price of Microsoft is lognormal tomorrow. Because 1 day is such a short period of time, this is almost indistinguishable from the assumption we do make—that the change in the stock price between today and tomorrow is normal.

to that for Microsoft shows that the standard deviation of the change in the value of the portfolio in 1 day is

$$5,000,000 \times 0.01 = 50,000$$

Assuming the change is normally distributed, the 1-day 99% VaR is

$$50,000 \times 2.33 = \$116,500$$

and the 10-day 99% VaR is

$$116,500 \times \sqrt{10} = \$368,405$$

Two-Asset Case

Now consider a portfolio consisting of both \$10 million of Microsoft shares and \$5 million of AT&T shares. We suppose that the returns on the two shares have a bivariate normal distribution with a correlation of 0.3. A standard result in statistics tells us that, if two variables X and Y have standard deviations equal to σ_X and σ_Y with the coefficient of correlation between them equal to ρ , the standard deviation of $X + Y$ is given by

$$\sigma_{X+Y} = \sqrt{\sigma_X^2 + \sigma_Y^2 + 2\rho\sigma_X\sigma_Y}$$

To apply this result, we set X equal to the change in the value of the position in Microsoft over a 1-day period and Y equal to the change in the value of the position in AT&T over a 1-day period, so that

$$\sigma_X = 200,000 \quad \text{and} \quad \sigma_Y = 50,000$$

The standard deviation of the change in the value of the portfolio consisting of both stocks over a 1-day period is therefore

$$\sqrt{200,000^2 + 50,000^2 + 2 \times 0.3 \times 200,000 \times 50,000} = 220,227$$

The mean change is assumed to be zero and the change is normally distributed. So the 1-day 99% VaR is therefore

$$220,227 \times 2.33 = \$513,129$$

The 10-day 99% VaR is $\sqrt{10}$ times this, or \$1,622,657.

The Benefits of Diversification

In the example we have just considered:

1. The 10-day 99% VaR for the portfolio of Microsoft shares is \$1,473,621.
2. The 10-day 99% VaR for the portfolio of AT&T shares is \$368,405.
3. The 10-day 99% VaR for the portfolio of both Microsoft and AT&T shares is \$1,622,657.

The amount

$$(1,473,621 + 368,405) - 1,622,657 = \$219,369$$

represents the benefits of diversification. If Microsoft and AT&T were perfectly correlated, the VaR for the portfolio of both Microsoft and AT&T would equal the VaR for the Microsoft portfolio plus the VaR for the AT&T portfolio. Less than perfect correlation leads to some of the risk being “diversified away.”⁵

21.4 THE LINEAR MODEL

The examples we have just considered are simple illustrations of the use of the linear model for calculating VaR. Suppose that we have a portfolio worth P consisting of n assets with an amount α_i being invested in asset i ($1 \leq i \leq n$). Define Δx_i as the return on asset i in one day. The dollar change in the value of our investment in asset i in one day is $\alpha_i \Delta x_i$ and

$$\Delta P = \sum_{i=1}^n \alpha_i \Delta x_i \quad (21.1)$$

where ΔP is the dollar change in the value of the whole portfolio in one day.

In the example considered in the previous section, \$10 million was invested in the first asset (Microsoft) and \$5 million was invested in the second asset (AT&T), so that (in millions of dollars) $\alpha_1 = 10$, $\alpha_2 = 5$, and

$$\Delta P = 10\Delta x_1 + 5\Delta x_2$$

If we assume that the Δx_i in equation (21.1) are multivariate normal, then ΔP is normally distributed. To calculate VaR, we therefore need to calculate only the mean and standard deviation of ΔP . We assume, as discussed in the previous section, that the expected value of each Δx_i is zero. This implies that the mean of ΔP is zero.

To calculate the standard deviation of ΔP , we define σ_i as the daily volatility of the i th asset and ρ_{ij} as the coefficient of correlation between returns on asset i and asset j . This means that σ_i is the standard deviation of Δx_i , and ρ_{ij} is the coefficient of correlation between Δx_i and Δx_j . The variance of ΔP , which we will denote by σ_P^2 , is given by

$$\sigma_P^2 = \sum_{i=1}^n \sum_{j=1}^n \rho_{ij} \alpha_i \alpha_j \sigma_i \sigma_j \quad (21.2)$$

This equation can also be written as

$$\sigma_P^2 = \sum_{i=1}^n \alpha_i^2 \sigma_i^2 + 2 \sum_{i=1}^n \sum_{j < i} \rho_{ij} \alpha_i \alpha_j \sigma_i \sigma_j$$

The standard deviation of the change over N days is $\sigma_P \sqrt{N}$, and the 99% VaR for an N -day time horizon is $2.33\sigma_P \sqrt{N}$.

⁵ Harry Markowitz was one of the first researchers to study the benefits of diversification to a portfolio manager. He was awarded a Nobel prize for this research in 1990. See H. Markowitz, “Portfolio Selection,” *Journal of Finance*, 7, 1 (March 1952): 77–91.

The portfolio return in one day is $\Delta P/P$. From equation (21.2), the variance of this is

$$\sum_{i=1}^n \sum_{j=1}^n \rho_{ij} w_i w_j \sigma_i \sigma_j$$

where $w_i = \alpha_i/P$ is the weight of the i th investment in the portfolio. This version of equation (21.2) is the one usually used by portfolio managers.

In the example considered in the previous section, $\sigma_1 = 0.02$, $\sigma_2 = 0.01$, and $\rho_{12} = 0.3$. As already noted, $\alpha_1 = 10$ and $\alpha_2 = 5$, so that

$$\sigma_p^2 = 10^2 \times 0.02^2 + 5^2 \times 0.01^2 + 2 \times 10 \times 5 \times 0.3 \times 0.02 \times 0.01 = 0.0485$$

and $\sigma_p = 0.220$. This is the standard deviation of the change in the portfolio value per day (in millions of dollars). The ten-day 99% VaR is $2.33 \times 0.220 \times \sqrt{10} = \1.623 million. This agrees with the calculation in the previous section.

Correlation and Covariance Matrices

A correlation matrix is a matrix where the entry in the i th row and j th column is the correlation ρ_{ij} between variable i and j . It is shown in Table 21.5. Since a variable is always perfectly correlated with itself, the diagonal elements of the correlation matrix are 1. Furthermore, because $\rho_{ij} = \rho_{ji}$, the correlation matrix is symmetric. The correlation matrix, together with the daily standard deviations of the variables, enables the portfolio variance to be calculated using equation (21.2).

Instead of working with correlations and volatilities, analysts often use variances and covariances. The daily variance var_i of variable i is the square of its daily volatility:

$$\text{var}_i = \sigma_i^2$$

The covariance cov_{ij} between variable i and variable j is the product of the daily volatility of variable i , the daily volatility of variable j , and the correlation between i and j :

$$\text{cov}_{ij} = \sigma_i \sigma_j \rho_{ij}$$

The equation for the variance of the portfolio in equation (21.2) can be written

$$\sigma_p^2 = \sum_{i=1}^n \sum_{j=1}^n \text{cov}_{ij} \alpha_i \alpha_j \quad (21.3)$$

Table 21.5 A correlation matrix: ρ_{ij} is the correlation between variable i and variable j .

$$\begin{bmatrix} 1 & \rho_{12} & \rho_{13} & \cdots & \rho_{1n} \\ \rho_{21} & 1 & \rho_{23} & \cdots & \rho_{2n} \\ \rho_{31} & \rho_{32} & 1 & \cdots & \rho_{3n} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \rho_{n1} & \rho_{n2} & \rho_{n3} & \cdots & 1 \end{bmatrix}$$

Table 21.6 A variance–covariance matrix: cov_{ij} is the covariance between variable i and variable j . Diagonal entries are variance: $\text{cov}_{ii} = \text{var}_i$

$$\begin{bmatrix} \text{var}_1 & \text{cov}_{12} & \text{cov}_{13} & \cdots & \text{cov}_{1n} \\ \text{cov}_{21} & \text{var}_2 & \text{cov}_{23} & \cdots & \text{cov}_{2n} \\ \text{cov}_{31} & \text{cov}_{32} & \text{var}_3 & \cdots & \text{cov}_{3n} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \text{cov}_{n1} & \text{cov}_{n2} & \text{cov}_{n3} & \cdots & \text{var}_n \end{bmatrix}$$

In a *covariance matrix*, the entry in the i th row and j th column is the covariance between variable i and variable j . As just mentioned, the covariance between a variable and itself is its variance. The diagonal entries in the matrix are therefore variances (see Table 21.6). For this reason, the covariance matrix is sometimes called the *variance–covariance matrix*. (Like the correlation matrix, it is symmetric.) Using matrix notation, the equation for the standard deviation of the portfolio just given becomes

$$\sigma_p^2 = \alpha C \alpha$$

where α is the (column) vector whose i th element is α_i , C is the variance–covariance matrix, and α is the transpose of α .

The variances and covariances are generally calculated from historical data. We will illustrate this in Section 22.8 for the four-index example introduced in Section 21.2.

Handling Interest Rates

It is out of the question in the model-building approach to define a separate market variable for every single bond price or interest rate to which a company is exposed. Some simplifications are necessary when the model-building approach is used. One possibility is to assume that only parallel shifts in the yield curve occur. It is then necessary to define only one market variable: the size of the parallel shift. The changes in the value of a bond portfolio can then be calculated using the duration relationship

$$\Delta P = -DP \Delta y$$

where P is the value of the portfolio, ΔP is the change in P in one day, D is the modified duration of the portfolio, and Δy is the parallel shift in 1 day.

This approach does not usually give enough accuracy. The procedure usually followed is to choose as market variables the prices of zero-coupon bonds with standard maturities: 1 month, 3 months, 6 months, 1 year, 2 years, 5 years, 7 years, 10 years, and 30 years. For the purposes of calculating VaR, the cash flows from instruments in the portfolio are mapped into cash flows occurring on the standard maturity dates. Consider a \$1 million position in a Treasury bond lasting 1.2 years that pays a coupon of 6% semiannually. Coupons are paid in 0.2, 0.7, and 1.2 years, and the principal is paid in 1.2 years. This bond is, therefore, in the first instance regarded as a \$30,000 position in 0.2-year zero-coupon bond plus a \$30,000 position in a 0.7-year zero-coupon bond plus a \$1.03 million position in a 1.2-year zero-coupon bond. The

position in the 0.2-year bond is then replaced by an equivalent position in 1-month and 3-month zero-coupon bonds; the position in the 0.7-year bond is replaced by an equivalent position in 6-month and 1-year zero-coupon bonds; and the position in the 1.2-year bond is replaced by an equivalent position in 1-year and 2-year zero-coupon bonds. The result is that the position in the 1.2-year coupon-bearing bond is for VaR purposes regarded as a position in zero-coupon bonds having maturities of 1 month, 3 months, 6 months, 1 year, and 2 years.

This procedure is known as *cash-flow mapping*. One way of doing it is explained in Technical Note 25 at www.rotman.utoronto.ca/~hull/TechnicalNotes. Note that cash-flow mapping is not necessary when the historical simulation approach is used. This is because the complete term structure of interest rates can be calculated for each of the scenarios considered.

Applications of the Linear Model

The simplest application of the linear model is to a portfolio with no derivatives consisting of positions in stocks, bonds, foreign exchange, and commodities. In this case, the change in the value of the portfolio is linearly dependent on the percentage changes in the prices of the assets comprising the portfolio. Note that, for the purposes of VaR calculations, all asset prices are measured in the domestic currency. The market variables considered by a large bank in the United States are therefore likely to include the value of the Nikkei 225 index measured in dollars, the price of a 10-year sterling zero-coupon bond measured in dollars, and so on.

An example of a derivative that can be handled by the linear model is a forward contract to buy a foreign currency. Suppose the contract matures at time T . It can be regarded as the exchange of a foreign zero-coupon bond maturing at time T for a domestic zero-coupon bond maturing at time T . For the purposes of calculating VaR, the forward contract is therefore treated as a long position in the foreign bond combined with a short position in the domestic bond. Each bond can be handled using a cash-flow mapping procedure.

Consider next an interest rate swap. As explained in Chapter 7, this can be regarded as the exchange of a floating-rate bond for a fixed-rate bond. The fixed-rate bond is a regular coupon-bearing bond. The floating-rate bond is worth par just after the next payment date. It can be regarded as a zero-coupon bond with a maturity date equal to the next payment date. The interest rate swap therefore reduces to a portfolio of long and short positions in bonds and can be handled using a cash-flow mapping procedure.

The Linear Model and Options

We now consider how we might try to use the linear model when there are options. Consider first a portfolio consisting of options on a single stock whose current price is S . Suppose that the delta of the position (calculated in the way described in Chapter 18) is δ .⁶ Since δ is the rate of change of the value of the portfolio with S , it is approximately true that

$$\delta = \frac{\Delta P}{\Delta S}$$

⁶ Normally we denote the delta and gamma of a portfolio by Δ and Γ . In this section and the next, we use the lower case Greek letters δ and γ to avoid overworking Δ .

or

$$\Delta P = \delta \Delta S \quad (21.4)$$

where ΔS is the dollar change in the stock price in 1 day and ΔP is, as usual, the dollar change in the portfolio in 1 day. Define Δx as the percentage change in the stock price in 1 day, so that

$$\Delta x = \frac{\Delta S}{S}$$

It follows that an approximate relationship between ΔP and Δx is

$$\Delta P = S\delta \Delta x$$

When we have a position in several underlying market variables that includes options, we can derive an approximate linear relationship between ΔP and the Δx_i similarly. This relationship is

$$\Delta P = \sum_{i=1}^n S_i \delta_i \Delta x_i \quad (21.5)$$

where S_i is the value of the i th market variable and δ_i is the delta of the portfolio with respect to the i th market variable. This corresponds to equation (21.1):

$$\Delta P = \sum_{i=1}^n \alpha_i \Delta x_i$$

with $\alpha_i = S_i \delta_i$. Equation (21.2) or (21.3) can therefore be used to calculate the standard deviation of ΔP .

Example 21.1

A portfolio consists of options on Microsoft and AT&T. The options on Microsoft have a delta of 1,000, and the options on AT&T have a delta of 20,000. The Microsoft share price is \$120, and the AT&T share price is \$30. From equation (21.5), it is approximately true that

$$\Delta P = 120 \times 1,000 \times \Delta x_1 + 30 \times 20,000 \times \Delta x_2$$

or

$$\Delta P = 120,000\Delta x_1 + 600,000\Delta x_2$$

where Δx_1 and Δx_2 are the returns from Microsoft and AT&T in 1 day and ΔP is the resultant change in the value of the portfolio. (The portfolio is assumed to be equivalent to an investment of \$120,000 in Microsoft and \$600,000 in AT&T.) Assuming that the daily volatility of Microsoft is 2% and the daily volatility of AT&T is 1% and the correlation between the daily changes is 0.3, the standard deviation of ΔP (in thousands of dollars) is

$$\sqrt{(120 \times 0.02)^2 + (600 \times 0.01)^2 + 2 \times 120 \times 0.02 \times 600 \times 0.01 \times 0.3} = 7.099$$

Since $N(-1.65) = 0.05$, the 5-day 95% VaR is $1.65 \times \sqrt{5} \times 7,099 = \$26,193$.

21.5 THE QUADRATIC MODEL

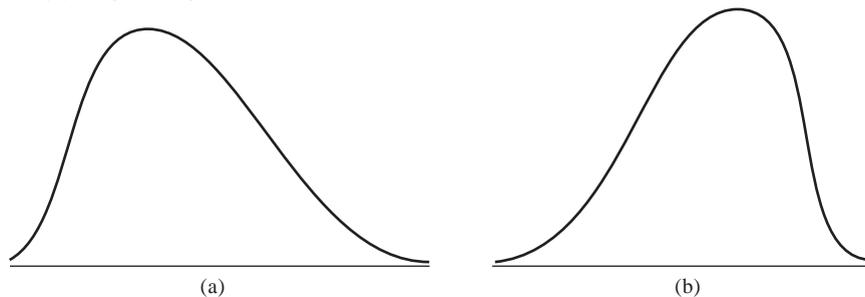
When a portfolio includes options, the linear model is an approximation. It does not take account of the gamma of the portfolio. As discussed in Chapter 18, delta is defined as the rate of change of the portfolio value with respect to an underlying market variable and gamma is defined as the rate of change of the delta with respect to the market variable. Gamma measures the curvature of the relationship between the portfolio value and an underlying market variable.

Figure 21.4 shows the impact of a nonzero gamma on the probability distribution of the value of the portfolio. When gamma is positive, the probability distribution tends to be positively skewed; when gamma is negative, it tends to be negatively skewed. Figures 21.5 and 21.6 illustrate the reason for this result. Figure 21.4 shows the relationship between the value of a long call option and the price of the underlying asset. A long call is an example of an option position with positive gamma. The figure shows that, when the probability distribution for the price of the underlying asset at the end of 1 day is normal, the probability distribution for the option price is positively skewed.⁷ Figure 21.6 shows the relationship between the value of a short call position and the price of the underlying asset. A short call position has a negative gamma. In this case, we see that a normal distribution for the price of the underlying asset at the end of 1 day gets mapped into a negatively skewed distribution for the value of the option position.

The VaR for a portfolio is critically dependent on the left tail of the probability distribution of the portfolio value. For example, when the confidence level used is 99%, the VaR is the value in the left tail below which there is only 1% of the distribution. As indicated in Figures 21.4a and 21.5, a positive gamma portfolio tends to have a less heavy left tail than the normal distribution. If the distribution of ΔP is normal, the calculated VaR tends to be too high. Similarly, as indicated in Figures 21.4b and 21.6, a negative gamma portfolio tends to have a heavier left tail than the normal distribution. If the distribution of ΔP is normal, the calculated VaR tends to be too low.

For a more accurate estimate of VaR than that given by the linear model, both delta and gamma measures can be used to relate ΔP to the Δx_i . Consider a portfolio dependent on a single asset whose price is S . Suppose δ and γ are the delta and gamma

Figure 21.4 Probability distribution for value of portfolio: (a) positive gamma; (b) negative gamma.



⁷ As mentioned in footnote 4, we can use the normal distribution as an approximation to the lognormal distribution in VaR calculations.

Figure 21.5 Translation of normal probability distribution for asset into probability distribution for value of a long call on asset.

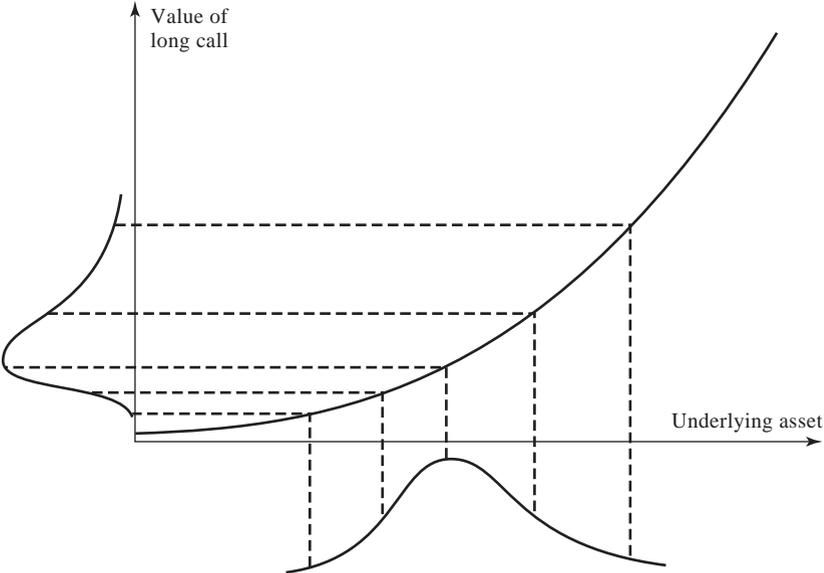
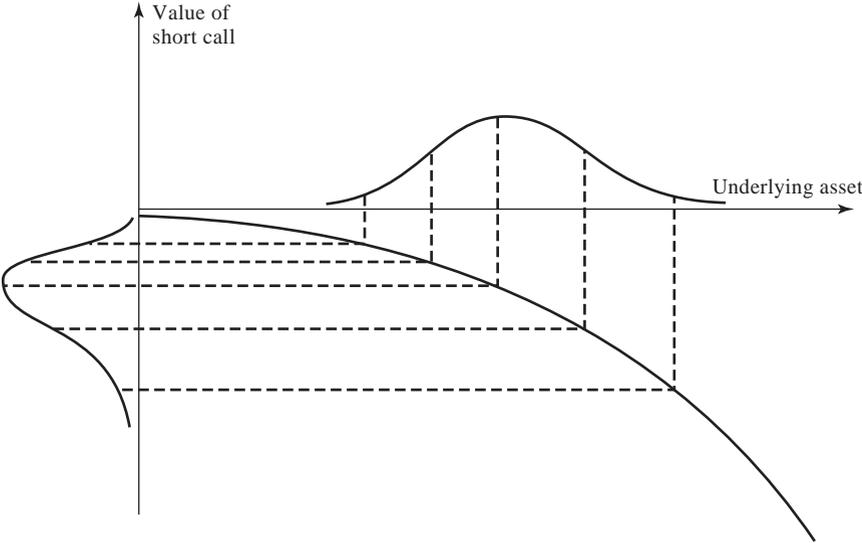


Figure 21.6 Translation of normal probability distribution for asset into probability distribution for value of a short call on asset.



of the portfolio. From the appendix to Chapter 18, the equation

$$\Delta P = \delta \Delta S + \frac{1}{2}\gamma(\Delta S)^2$$

is an improvement over the approximation in equation (21.4).⁸ Setting

$$\Delta x = \frac{\Delta S}{S}$$

reduces this to

$$\Delta P = S\delta \Delta x + \frac{1}{2}S^2\gamma(\Delta x)^2 \quad (21.6)$$

More generally for a portfolio with n underlying market variables, with each instrument in the portfolio being dependent on only one of the market variables, equation (21.6) becomes

$$\Delta P = \sum_{i=1}^n S_i \delta_i \Delta x_i + \sum_{i=1}^n \frac{1}{2} S_i^2 \gamma_i (\Delta x_i)^2$$

where S_i is the value of the i th market variable, and δ_i and γ_i are the delta and gamma of the portfolio with respect to the i th market variable. When individual instruments in the portfolio may be dependent on more than one market variable, this equation takes the more general form

$$\Delta P = \sum_{i=1}^n S_i \delta_i \Delta x_i + \sum_{i=1}^n \sum_{j=1}^n \frac{1}{2} S_i S_j \gamma_{ij} \Delta x_i \Delta x_j \quad (21.7)$$

where γ_{ij} is a “cross gamma” defined as

$$\gamma_{ij} = \frac{\partial^2 P}{\partial S_i \partial S_j}$$

Equation (21.7) is not as easy to work with as equation (21.1), but it can be used to calculate moments for ΔP . A result in statistics known as the Cornish–Fisher expansion can be used to estimate percentiles of the probability distribution from the moments.⁹

21.6 MONTE CARLO SIMULATION

As an alternative to the procedure described so far, the model-building approach can be implemented using Monte Carlo simulation to generate the probability distribution

⁸ The Taylor series expansion in the appendix to Chapter 18 suggests the approximation

$$\Delta P = \Theta \Delta t + \delta \Delta S + \frac{1}{2}\gamma(\Delta S)^2$$

when terms of higher order than Δt are ignored. In practice, the $\Theta \Delta t$ term is so small that it is usually ignored.

⁹ See Technical Note 10 at www.rotman.utoronto.ca/~hull/TechnicalNotes for details of the calculation of moments and the use of Cornish–Fisher expansions. When there is a single underlying variable, $E(\Delta P) = 0.5S^2\gamma\sigma^2$, $E(\Delta P^2) = S^2\delta^2\sigma^2 + 0.75S^4\gamma^2\sigma^4$, and $E(\Delta P^3) = 4.5S^4\delta^2\gamma\sigma^4 + 1.875S^6\gamma^3\sigma^6$, where S is the value of the variable and σ is its daily volatility. Sample Application E in the DerivaGem Applications implements the Cornish–Fisher expansion method for this case.

for ΔP . Suppose we wish to calculate a 1-day VaR for a portfolio. The procedure is as follows:

1. Value the portfolio today in the usual way using the current values of market variables.
2. Sample once from the multivariate normal probability distribution of the Δx_i .¹⁰
3. Use the values of the Δx_i that are sampled to determine the value of each market variable at the end of one day.
4. Revalue the portfolio at the end of the day in the usual way.
5. Subtract the value calculated in Step 1 from the value in Step 4 to determine a sample ΔP .
6. Repeat Steps 2 to 5 many times to build up a probability distribution for ΔP .

The VaR is calculated as the appropriate percentile of the probability distribution of ΔP . Suppose, for example, that we calculate 5,000 different sample values of ΔP in the way just described. The 1-day 99% VaR is the value of ΔP for the 50th worst outcome; the 1-day VaR 95% is the value of ΔP for the 250th worst outcome; and so on.¹¹ The N -day VaR is usually assumed to be the 1-day VaR multiplied by \sqrt{N} .¹²

The drawback of Monte Carlo simulation is that it tends to be slow because a company's complete portfolio (which might consist of hundreds of thousands of different instruments) has to be revalued many times.¹³ One way of speeding things up is to assume that equation (21.7) describes the relationship between ΔP and the Δx_i . We can then jump straight from Step 2 to Step 5 in the Monte Carlo simulation and avoid the need for a complete revaluation of the portfolio. This is sometimes referred to as the *partial simulation approach*. A similar approach is sometimes used when implementing historical simulation.

21.7 COMPARISON OF APPROACHES

We have discussed two methods for estimating VaR: the historical simulation approach and the model-building approach. The advantages of the model-building approach are that results can be produced very quickly and it can easily be used in conjunction with volatility updating schemes such as those we will describe in the next chapter. The main disadvantage of the model-building approach is that it assumes that the market variables have a multivariate normal distribution. In practice, daily changes in market variables often have distributions with tails that are quite different from the normal distribution. This is illustrated in Table 19.1.

The historical simulation approach has the advantage that historical data determine the joint probability distribution of the market variables. It also avoids the need for

¹⁰ One way of doing so is given in Section 20.6.

¹¹ As in the case of historical simulation, extreme value theory can be used to "smooth the tails" so that better estimates of extreme percentiles are obtained.

¹² This is only approximately true when the portfolio includes options, but it is the assumption that is made in practice for most VaR calculation methods.

¹³ An approach for limiting the number of portfolio revaluations is proposed in F. Jamshidian and Y. Zhu "Scenario simulation model: theory and methodology," *Finance and Stochastics*, 1 (1997), 43–67.

cash-flow mapping. The main disadvantages of historical simulation are that it is computationally slow and does not easily allow volatility updating schemes to be used.¹⁴

One disadvantage of the model-building approach is that it tends to give poor results for low-delta portfolios (see Problem 21.21).

21.8 STRESS TESTING AND BACK TESTING

In addition to calculating VaR, many companies carry out what is known as *stress testing*. This involves estimating how a company's portfolio would have performed under some of the most extreme market moves seen in the last 10 to 20 years.

For example, to test the impact of an extreme movement in US equity prices, a company might set the percentage changes in all market variables equal to those on October 19, 1987 (when the S&P 500 moved by 22.3 standard deviations). If this is considered to be too extreme, the company might choose January 8, 1988 (when the S&P 500 moved by 6.8 standard deviations). To test the effect of extreme movements in UK interest rates, the company might set the percentage changes in all market variables equal to those on April 10, 1992 (when 10-year bond yields moved by 7.7 standard deviations).

The scenarios used in stress testing are also sometimes generated by senior management. One technique sometimes used is to ask senior management to meet periodically and “brainstorm” to develop extreme scenarios that might occur given the current economic environment and global uncertainties.

Stress testing can be considered as a way of taking into account extreme events that do occur from time to time but are virtually impossible according to the probability distributions assumed for market variables. A 5-standard-deviation daily move in a market variable is one such extreme event. Under the assumption of a normal distribution, it happens about once every 7,000 years, but, in practice, it is not uncommon to see a 5-standard-deviation daily move once or twice every 10 years.

Following the credit crisis of 2007 and 2008, regulators have proposed the calculation of *stressed VaR*. This is VaR based on a historical simulation of how market variables moved during a period of stressed market conditions (such as those in 2008).

Whatever the method used for calculating VaR, an important reality check is *back testing*. It involves testing how well the VaR estimates would have performed in the past. Suppose that we are calculating a 1-day 99% VaR. Back testing would involve looking at how often the loss in a day exceeded the 1-day 99% VaR that would have been calculated for that day. If this happened on about 1% of the days, we can feel reasonably comfortable with the methodology for calculating VaR. If it happened on, say, 7% of days, the methodology is suspect.

21.9 PRINCIPAL COMPONENTS ANALYSIS

One approach to handling the risk arising from groups of highly correlated market variables is principal components analysis. This takes historical data on movements in

¹⁴ For a way of adapting the historical simulation approach to incorporate volatility updating, see J. Hull and A. White. “Incorporating volatility updating into the historical simulation method for value-at-risk,” *Journal of Risk* 1, No. 1 (1998): 5–19.

Table 21.7 Factor loadings for US Treasury data.

	<i>PC1</i>	<i>PC2</i>	<i>PC3</i>	<i>PC4</i>	<i>PC5</i>	<i>PC6</i>	<i>PC7</i>	<i>PC8</i>	<i>PC9</i>	<i>PC10</i>
3m	0.21	-0.57	0.50	0.47	-0.39	-0.02	0.01	0.00	0.01	0.00
6m	0.26	-0.49	0.23	-0.37	0.70	0.01	-0.04	-0.02	-0.01	0.00
12m	0.32	-0.32	-0.37	-0.58	-0.52	-0.23	-0.04	-0.05	0.00	0.01
2y	0.35	-0.10	-0.38	0.17	0.04	0.59	0.56	0.12	-0.12	-0.05
3y	0.36	0.02	-0.30	0.27	0.07	0.24	-0.79	0.00	-0.09	-0.00
4y	0.36	0.14	-0.12	0.25	0.16	-0.63	0.15	0.55	-0.14	-0.08
5y	0.36	0.17	-0.04	0.14	0.08	-0.10	0.09	-0.26	0.71	0.48
7y	0.34	0.27	0.15	0.01	0.00	-0.12	0.13	-0.54	0.00	-0.68
10y	0.31	0.30	0.28	-0.10	-0.06	0.01	0.03	-0.23	-0.63	0.52
30y	0.25	0.33	0.46	-0.34	-0.18	0.33	-0.09	0.52	0.26	-0.13

the market variables and attempts to define a set of components or factors that explain the movements.

The approach is best illustrated with an example. The market variables we will consider are 10 US Treasury rates with maturities between 3 months and 30 years. Tables 21.7 and 21.8 shows results produced by Frye for these market variables using 1,543 daily observations between 1989 and 1995.¹⁵ The first column in Table 21.7 shows the maturities of the rates that were considered. The remaining 10 columns in the table show the 10 factors (or principal components) describing the rate moves. The first factor, shown in the column labeled PC1, corresponds to a roughly parallel shift in the yield curve. When there is one unit of that factor, the 3-month rate increases by 0.21 basis points, the 6-month rate increases by 0.26 basis points, and so on. The second factor is shown in the column labeled PC2. It corresponds to a “twist” or “steepening” of the yield curve. Rates between 3 months and 2 years move in one direction; rates between 3 years and 30 years move in the other direction. The third factor corresponds to a “bowing” of the yield curve. Rates at the short end and long end of the yield curve move in one direction; rates in the middle move in the other direction. The interest rate move for a particular factor is known as *factor loading*. In the example, the first factor’s loading for the three-month rate is 0.21.¹⁶

Because there are 10 rates and 10 factors, the interest rate changes observed on any given day can always be expressed as a linear sum of the factors by solving a set of 10 simultaneous equations. The quantity of a particular factor in the interest rate changes on a particular day is known as the *factor score* for that day.

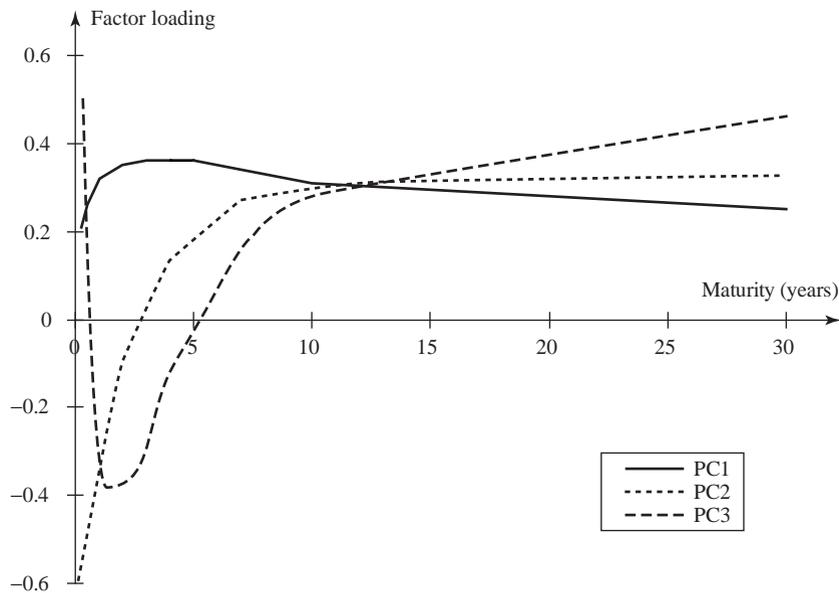
Table 21.8 Standard deviation of factor scores.

<i>PC1</i>	<i>PC2</i>	<i>PC3</i>	<i>PC4</i>	<i>PC5</i>	<i>PC6</i>	<i>PC7</i>	<i>PC8</i>	<i>PC9</i>	<i>PC10</i>
17.49	6.05	3.10	2.17	1.97	1.69	1.27	1.24	0.80	0.79

¹⁵ See J. Frye, “Principals of Risk: Finding VAR through Factor-Based Interest Rate Scenarios,” in *VAR: Understanding and Applying Value at Risk*, pp. 275–88. London: Risk Publications, 1997.

¹⁶ The factor loadings have the property that the sum of their squares for each factor is 1.0.

Figure 21.7 The three most important factors driving yield curve movements.



The importance of a factor is measured by the standard deviation of its factor score. The standard deviations of the factor scores in our example are shown in Table 21.8 and the factors are listed in order of their importance. The numbers in Table 21.8 are measured in basis points. A quantity of the first factor equal to one standard deviation, therefore, corresponds to the 3-month rate moving by $0.21 \times 17.49 = 3.67$ basis points, the 6-month rate moving by $0.26 \times 17.49 = 4.55$ basis points, and so on.

The technical details of how the factors are determined are not covered here. It is sufficient for us to note that the factors are chosen so that the factor scores are uncorrelated. For instance, in our example, the first factor score (amount of parallel shift) is uncorrelated with the second factor score (amount of twist) across the 1,543 days. The variances of the factor scores (i.e., the squares of the standard deviations) have the property that they add up to the total variance of the data. From Table 21.8, the total variance of the original data (i.e., sum of the variance of the observations on the 3-month rate, the variance of the observations on the 6-month rate, and so on) is

$$17.49^2 + 6.05^2 + 3.10^2 + \dots + 0.79^2 = 367.9$$

From this it can be seen that the first factor accounts for $17.49^2/367.9 = 83.1\%$ of the variance in the original data; the first two factors account for $(17.49^2 + 6.05^2)/367.9 = 93.1\%$ of the variance in the data; the third factor accounts for a further 2.8% of the variance. This shows most of the risk in interest rate moves is accounted for by the first two or three factors. It suggests that we can relate the risks in a portfolio of interest rate dependent instruments to movements in these factors instead of considering all

ten interest rates. The three most important factors from Table 21.7 are plotted in Figure 21.7.¹⁷

Using Principal Components Analysis to Calculate VaR

To illustrate how a principal components analysis can be used to calculate VaR, consider a portfolio with the exposures to interest rate moves shown in Table 21.9. A 1-basis-point change in the 1-year rate causes the portfolio value to increase by \$10 million, a 1-basis-point change in the 2-year rate causes it to increase by \$4 million, and so on. Suppose the first two factors are used to model rate moves. (As mentioned above, this captures 93.1% of the variance in rate moves.) Using the data in Table 21.7, the exposure to the first factor (measured in millions of dollars per factor score basis point) is

$$10 \times 0.32 + 4 \times 0.35 - 8 \times 0.36 - 7 \times 0.36 + 2 \times 0.36 = -0.08$$

and the exposure to the second factor is

$$10 \times (-0.32) + 4 \times (-0.10) - 8 \times 0.02 - 7 \times 0.14 + 2 \times 0.17 = -4.40$$

Suppose that f_1 and f_2 are the factor scores (measured in basis points). The change in the portfolio value is, to a good approximation, given by

$$\Delta P = -0.08f_1 - 4.40f_2$$

The factor scores are uncorrelated and have the standard deviations given in Table 21.8. The standard deviation of ΔP is therefore

$$\sqrt{0.08^2 \times 17.49^2 + 4.40^2 \times 6.05^2} = 26.66$$

Hence, the 1-day 99% VaR is $26.66 \times 2.33 = 62.12$. Note that the data in Table 21.9 are such that there is very little exposure to the first factor and significant exposure to the second factor. Using only one factor would significantly understate VaR (see Problem 21.11). The duration-based method for handling interest rates, mentioned in Section 21.4, would also significantly understate VaR as it considers only parallel shifts in the yield curve.

A principal components analysis can in theory be used for market variables other than interest rates. Suppose that a financial institution has exposures to a number of different stock indices. A principal components analysis can be used to identify

Table 21.9 Change in portfolio value for a 1-basis-point rate move (\$ millions).

<i>1-year rate</i>	<i>2-year rate</i>	<i>3-year rate</i>	<i>4-year rate</i>	<i>5-year rate</i>
+10	+4	-8	-7	+2

¹⁷ Similar results to those described here, in respect of the nature of the factors and the amount of the total risk they account for, are obtained when a principal components analysis is used to explain the movements in almost any yield curve in any country.

factors describing movements in the indices and the most important of these can be used to replace the market indices in a VaR analysis. How effective a principal components analysis is for a group of market variables depends on how closely correlated they are.

As explained earlier in the chapter, VaR is usually calculated by relating the actual changes in a portfolio to percentage changes in market variables (the Δx_i). For a VaR calculation, it may therefore be most appropriate to carry out a principal components analysis on percentage changes in market variables rather than actual changes.

SUMMARY

A value at risk (VaR) calculation is aimed at making a statement of the form: “We are X percent certain that we will not lose more than V dollars in the next N days.” The variable V is the VaR, $X\%$ is the confidence level, and N days is the time horizon.

One approach to calculating VaR is historical simulation. This involves creating a database consisting of the daily movements in all market variables over a period of time. The first simulation trial assumes that the percentage changes in each market variable are the same as those on the first day covered by the database; the second simulation trial assumes that the percentage changes are the same as those on the second day; and so on. The change in the portfolio value, ΔP , is calculated for each simulation trial, and the VaR is calculated as the appropriate percentile of the probability distribution of ΔP .

An alternative is the model-building approach. This is relatively straightforward if two assumptions can be made:

1. The change in the value of the portfolio (ΔP) is linearly dependent on percentage changes in market variables.
2. The percentage changes in market variables are multivariate normally distributed.

The probability distribution of ΔP is then normal, and there are analytic formulas for relating the standard deviation of ΔP to the volatilities and correlations of the underlying market variables. The VaR can be calculated from well-known properties of the normal distribution.

When a portfolio includes options, ΔP is not linearly related to the percentage changes in market variables. From knowledge of the gamma of the portfolio, we can derive an approximate quadratic relationship between ΔP and percentage changes in market variables. Monte Carlo simulation can then be used to estimate VaR.

In the next chapter we discuss how volatilities and correlations can be estimated and monitored.

FURTHER READING

Artzner P., F. Delbaen, J.-M. Eber, and D. Heath. “Coherent Measures of Risk,” *Mathematical Finance*, 9 (1999): 203–28.

Basak, S., and A. Shapiro. “Value-at-Risk-Based Risk Management: Optimal Policies and Asset Prices,” *Review of Financial Studies*, 14, 2 (2001): 371–405.

- Beder, T. "VaR: Seductive but Dangerous," *Financial Analysts Journal*, 51, 5 (1995): 12–24.
- Boudoukh, J., M. Richardson, and R. Whitelaw. "The Best of Both Worlds," *Risk*, May 1998: 64–67.
- Dowd, K. *Beyond Value at Risk: The New Science of Risk Management*. New York: Wiley, 1998.
- Duffie, D., and J. Pan. "An Overview of Value at Risk," *Journal of Derivatives*, 4, 3 (Spring 1997): 7–49.
- Embrechts, P., C. Kluppelberg, and T. Mikosch. *Modeling Extremal Events for Insurance and Finance*. New York: Springer, 1997.
- Frye, J. "Principals of Risk: Finding VAR through Factor-Based Interest Rate Scenarios" in *VAR: Understanding and Applying Value at Risk*, pp. 275–88. London: Risk Publications, 1997.
- Hendricks, D. "Evaluation of Value-at-Risk Models Using Historical Data," *Economic Policy Review*, Federal Reserve Bank of New York, 2 (April 1996): 39–69.
- Hopper, G. "Value at Risk: A New Methodology for Measuring Portfolio Risk," *Business Review*, Federal Reserve Bank of Philadelphia, July/August 1996: 19–29.
- Hua P., and P. Wilmott, "Crash Courses," *Risk*, June 1997: 64–67.
- Hull, J. C., and A. White. "Value at Risk When Daily Changes in Market Variables Are Not Normally Distributed," *Journal of Derivatives*, 5 (Spring 1998): 9–19.
- Hull, J. C., and A. White. "Incorporating Volatility Updating into the Historical Simulation Method for Value at Risk," *Journal of Risk*, 1, 1 (1998): 5–19.
- Jackson, P., D. J. Maude, and W. Perraudin. "Bank Capital and Value at Risk," *Journal of Derivatives*, 4, 3 (Spring 1997): 73–90.
- Jamshidian, F., and Y. Zhu. "Scenario Simulation Model: Theory and Methodology," *Finance and Stochastics*, 1 (1997): 43–67.
- Jorion, P. *Value at Risk*, 3rd edn. McGraw-Hill, 2007.
- Longin, F. M. "Beyond the VaR," *Journal of Derivatives*, 8, 4 (Summer 2001): 36–48.
- Marshall, C., and M. Siegel. "Value at Risk: Implementing a Risk Measurement Standard," *Journal of Derivatives* 4, 3 (Spring 1997): 91–111.
- McNeil, A. J. "Extreme Value Theory for Risk Managers," in *Internal Modeling and CAD II*, London: Risk Books, 1999. See also: www.math.ethz.ch/~mcneil.
- Neftci, S. "Value at Risk Calculations, Extreme Events and Tail Estimation," *Journal of Derivatives*, 7, 3 (Spring 2000): 23–38.
- Rich, D. "Second Generation VaR and Risk-Adjusted Return on Capital," *Journal of Derivatives*, 10, 4 (Summer 2003): 51–61.

Practice Questions (Answers in Solutions Manual)

- 21.1. Consider a position consisting of a \$100,000 investment in asset A and a \$100,000 investment in asset B. Assume that the daily volatilities of both assets are 1% and that the coefficient of correlation between their returns is 0.3. What is the 5-day 99% VaR for the portfolio?
- 21.2. Describe three ways of handling instruments that are dependent on interest rates when the model-building approach is used to calculate VaR. How would you handle these instruments when historical simulation is used to calculate VaR?
- 21.3. A financial institution owns a portfolio of options on the US dollar–sterling exchange rate. The delta of the portfolio is 56.0. The current exchange rate is 1.5000. Derive an approximate linear relationship between the change in the portfolio value and the

- percentage change in the exchange rate. If the daily volatility of the exchange rate is 0.7%, estimate the 10-day 99% VaR.
- 21.4. Suppose you know that the gamma of the portfolio in the previous question is 16.2. How does this change your estimate of the relationship between the change in the portfolio value and the percentage change in the exchange rate?
 - 21.5. Suppose that the daily change in the value of a portfolio is, to a good approximation, linearly dependent on two factors, calculated from a principal components analysis. The delta of a portfolio with respect to the first factor is 6 and the delta with respect to the second factor is -4 . The standard deviations of the factor are 20 and 8, respectively. What is the 5-day 90% VaR?
 - 21.6. Suppose that a company has a portfolio consisting of positions in stocks, bonds, foreign exchange, and commodities. Assume that there are no derivatives. Explain the assumptions underlying (a) the linear model and (b) the historical simulation model for calculating VaR.
 - 21.7. Explain how an interest rate swap is mapped into a portfolio of zero-coupon bonds with standard maturities for the purposes of a VaR calculation.
 - 21.8. Explain the difference between value at risk and expected shortfall.
 - 21.9. Explain why the linear model can provide only approximate estimates of VaR for a portfolio containing options.
 - 21.10. Some time ago a company entered into a forward contract to buy £1 million for \$1.5 million. The contract now has 6 months to maturity. The daily volatility of a 6-month zero-coupon sterling bond (when its price is translated to dollars) is 0.06% and the daily volatility of a 6-month zero-coupon dollar bond is 0.05%. The correlation between returns from the two bonds is 0.8. The current exchange rate is 1.53. Calculate the standard deviation of the change in the dollar value of the forward contract in 1 day. What is the 10-day 99% VaR? Assume that the 6-month interest rate in both sterling and dollars is 5% per annum with continuous compounding.
 - 21.11. The text calculates a VaR estimate for the example in Table 21.9 assuming two factors. How does the estimate change if you assume (a) one factor and (b) three factors.
 - 21.12. A bank has a portfolio of options on an asset. The delta of the options is -30 and the gamma is -5 . Explain how these numbers can be interpreted. The asset price is 20 and its volatility is 1% per day. Adapt Sample Application E in the DerivaGem Application Builder software to calculate VaR.
 - 21.13. Suppose that in Problem 21.12 the vega of the portfolio is -2 per 1% change in the annual volatility. Derive a model relating the change in the portfolio value in 1 day to delta, gamma, and vega. Explain without doing detailed calculations how you would use the model to calculate a VaR estimate.
 - 21.14. The one-day 99% VaR is calculated for the four-index example in Section 21.2 as \$253,385. Look at the underlying spreadsheets on the author's website and calculate: (a) the one-day 95% VaR and (b) the one-day 97% VaR.
 - 21.15. Use the spreadsheets on the author's website to calculate the one-day 99% VaR, using the basic methodology in Section 21.2, if the four-index portfolio considered in Section 21.2 is equally divided between the four indices.

Further Questions

- 21.16. A company has a position in bonds worth \$6 million. The modified duration of the portfolio is 5.2 years. Assume that only parallel shifts in the yield curve can take place and that the standard deviation of the daily yield change (when yield is measured in percent) is 0.09. Use the duration model to estimate the 20-day 90% VaR for the portfolio. Explain carefully the weaknesses of this approach to calculating VaR. Explain two alternatives that give more accuracy.
- 21.17. Consider a position consisting of a \$300,000 investment in gold and a \$500,000 investment in silver. Suppose that the daily volatilities of these two assets are 1.8% and 1.2%, respectively, and that the coefficient of correlation between their returns is 0.6. What is the 10-day 97.5% VaR for the portfolio? By how much does diversification reduce the VaR?
- 21.18. Consider a portfolio of options on a single asset. Suppose that the delta of the portfolio is 12, the value of the asset is \$10, and the daily volatility of the asset is 2%. Estimate the 1-day 95% VaR for the portfolio from the delta. Suppose next that the gamma of the portfolio is -2.6 . Derive a quadratic relationship between the change in the portfolio value and the percentage change in the underlying asset price in one day. How would you use this in a Monte Carlo simulation?
- 21.19. A company has a long position in a 2-year bond and a 3-year bond, as well as a short position in a 5-year bond. Each bond has a principal of \$100 and pays a 5% coupon annually. Calculate the company's exposure to the 1-year, 2-year, 3-year, 4-year, and 5-year rates. Use the data in Tables 21.7 and 21.8 to calculate a 20-day 95% VaR on the assumption that rate changes are explained by (a) one factor, (b) two factors, and (c) three factors. Assume that the zero-coupon yield curve is flat at 5%.
- 21.20. A bank has written a call option on one stock and a put option on another stock. For the first option the stock price is 50, the strike price is 51, the volatility is 28% per annum, and the time to maturity is 9 months. For the second option the stock price is 20, the strike price is 19, the volatility is 25% per annum, and the time to maturity is 1 year. Neither stock pays a dividend, the risk-free rate is 6% per annum, and the correlation between stock price returns is 0.4. Calculate a 10-day 99% VaR:
- (a) Using only deltas
 - (b) Using the partial simulation approach
 - (c) Using the full simulation approach.
- 21.21. A common complaint of risk managers is that the model-building approach (either linear or quadratic) does not work well when delta is close to zero. Test what happens when delta is close to zero by using Sample Application E in the DerivaGem Applications. (You can do this by experimenting with different option positions and adjusting the position in the underlying to give a delta of zero.) Explain the results you get.
- 21.22. Suppose that the portfolio considered in Section 20.2 has (in \$000s) 3,000 in DJIA, 3,000 in FTSE, 1,000 in CAC 40 and 3,000 in Nikkei 225. Use the spreadsheet on the author's website to calculate what difference this makes to the one-day 99% VaR that is calculated in Section 21.2.



CHAPTER 22

Estimating Volatilities and Correlations

In this chapter we explain how historical data can be used to produce estimates of the current and future levels of volatilities and correlations. The chapter is relevant both to the calculation of value at risk using the model-building approach and to the valuation of derivatives. When calculating value at risk, we are most interested in the current levels of volatilities and correlations because we are assessing possible changes in the value of a portfolio over a very short period of time. When valuing derivatives, forecasts of volatilities and correlations over the whole life of the derivative are usually required.

The chapter considers models with imposing names such as exponentially weighted moving average (EWMA), autoregressive conditional heteroscedasticity (ARCH), and generalized autoregressive conditional heteroscedasticity (GARCH). The distinctive feature of the models is that they recognize that volatilities and correlations are not constant. During some periods, a particular volatility or correlation may be relatively low, whereas during other periods it may be relatively high. The models attempt to keep track of the variations in the volatility or correlation through time.

22.1 ESTIMATING VOLATILITY

Define σ_n as the volatility of a market variable on day n , as estimated at the end of day $n - 1$. The square of the volatility, σ_n^2 , on day n is the *variance rate*. We described the standard approach to estimating σ_n from historical data in Section 14.4. Suppose that the value of the market variable at the end of day i is S_i . The variable u_i is defined as the continuously compounded return during day i (between the end of day $i - 1$ and the end of day i):

$$u_i = \ln \frac{S_i}{S_{i-1}}$$

An unbiased estimate of the variance rate per day, σ_n^2 , using the most recent m observations on the u_i is

$$\sigma_n^2 = \frac{1}{m-1} \sum_{i=1}^m (u_{n-i} - \bar{u})^2 \quad (22.1)$$

where \bar{u} is the mean of the u_i s:

$$\bar{u} = \frac{1}{m} \sum_{i=1}^m u_{n-i}$$

For the purposes of monitoring daily volatility, the formula in equation (22.1) is usually changed in a number of ways:

1. u_i is defined as the percentage change in the market variable between the end of day $i - 1$ and the end of day i , so that:¹

$$u_i = \frac{S_i - S_{i-1}}{S_{i-1}} \quad (22.2)$$

2. \bar{u} is assumed to be zero.²

3. $m - 1$ is replaced by m .³

These three changes make very little difference to the estimates that are calculated, but they allow us to simplify the formula for the variance rate to

$$\sigma_n^2 = \frac{1}{m} \sum_{i=1}^m u_{n-i}^2 \quad (22.3)$$

where u_i is given by equation (22.2).⁴

Weighting Schemes

Equation (22.3) gives equal weight to $u_{n-1}^2, u_{n-2}^2, \dots, u_{n-m}^2$. Our objective is to estimate the current level of volatility, σ_n . It therefore makes sense to give more weight to recent data. A model that does this is

$$\sigma_n^2 = \sum_{i=1}^m \alpha_i u_{n-i}^2 \quad (22.4)$$

The variable α_i is the amount of weight given to the observation i days ago. The α 's are positive. If we choose them so that $\alpha_i < \alpha_j$ when $i > j$, less weight is given to older observations. The weights must sum to unity, so that

$$\sum_{i=1}^m \alpha_i = 1$$

¹ This is consistent with the point made in Section 21.3 about the way that volatility is defined for the purposes of VaR calculations.

² As explained in Section 21.3, this assumption usually has very little effect on estimates of the variance because the expected change in a variable in one day is very small when compared with the standard deviation of changes.

³ Replacing $m - 1$ by m moves us from an unbiased estimate of the variance to a maximum likelihood estimate. Maximum likelihood estimates are discussed later in the chapter.

⁴ Note that the u 's in this chapter play the same role as the Δx 's in Chapter 21. Both are daily percentage changes in market variables. In the case of the u 's, the subscripts count observations made on different days on the same market variable. In the case of the Δx 's, they count observations made on the same day on different market variables. The use of subscripts for σ is similarly different between the two chapters. In this chapter, the subscripts refer to days; in Chapter 21 they referred to market variables.

An extension of the idea in equation (22.4) is to assume that there is a long-run average variance rate and that this should be given some weight. This leads to the model that takes the form

$$\sigma_n^2 = \gamma V_L + \sum_{i=1}^m \alpha_i u_{n-i}^2 \quad (22.5)$$

where V_L is the long-run variance rate and γ is the weight assigned to V_L . Since the weights must sum to unity, it follows that

$$\gamma + \sum_{i=1}^m \alpha_i = 1$$

This is known as an ARCH(m) model. It was first suggested by Engle.⁵ The estimate of the variance is based on a long-run average variance and m observations. The older an observation, the less weight it is given. Defining $\omega = \gamma V_L$, the model in equation (22.5) can be written

$$\sigma_n^2 = \omega + \sum_{i=1}^m \alpha_i u_{n-i}^2 \quad (22.6)$$

In the next two sections we discuss two important approaches to monitoring volatility using the ideas in equations (22.4) and (22.5).

22.2 THE EXPONENTIALLY WEIGHTED MOVING AVERAGE MODEL

The exponentially weighted moving average (EWMA) model is a particular case of the model in equation (22.4) where the weights α_i decrease exponentially as we move back through time. Specifically, $\alpha_{i+1} = \lambda \alpha_i$, where λ is a constant between 0 and 1.

It turns out that this weighting scheme leads to a particularly simple formula for updating volatility estimates. The formula is

$$\sigma_n^2 = \lambda \sigma_{n-1}^2 + (1 - \lambda) u_{n-1}^2 \quad (22.7)$$

The estimate, σ_n , of the volatility of a variable for day n (made at the end of day $n - 1$) is calculated from σ_{n-1} (the estimate that was made at the end of day $n - 2$ of the volatility for day $n - 1$) and u_{n-1} (the most recent daily percentage change in the variable).

To understand why equation (22.7) corresponds to weights that decrease exponentially, we substitute for σ_{n-1}^2 to get

$$\sigma_n^2 = \lambda[\lambda \sigma_{n-2}^2 + (1 - \lambda) u_{n-2}^2] + (1 - \lambda) u_{n-1}^2$$

or

$$\sigma_n^2 = (1 - \lambda)(u_{n-1}^2 + \lambda u_{n-2}^2) + \lambda^2 \sigma_{n-2}^2$$

Substituting in a similar way for σ_{n-2}^2 gives

$$\sigma_n^2 = (1 - \lambda)(u_{n-1}^2 + \lambda u_{n-2}^2 + \lambda^2 u_{n-3}^2) + \lambda^3 \sigma_{n-3}^2$$

⁵ See R. Engle "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of UK Inflation," *Econometrica*, 50 (1982): 987–1008.

Continuing in this way gives

$$\sigma_n^2 = (1 - \lambda) \sum_{i=1}^m \lambda^{i-1} u_{n-i}^2 + \lambda^m \sigma_{n-m}^2$$

For large m , the term $\lambda^m \sigma_{n-m}^2$ is sufficiently small to be ignored, so that equation (22.7) is the same as equation (22.4) with $\alpha_i = (1 - \lambda)\lambda^{i-1}$. The weights for the u_i decline at rate λ as we move back through time. Each weight is λ times the previous weight.

Example 22.1

Suppose that λ is 0.90, the volatility estimated for a market variable for day $n - 1$ is 1% per day, and during day $n - 1$ the market variable increased by 2%. This means that $\sigma_{n-1}^2 = 0.01^2 = 0.0001$ and $u_{n-1}^2 = 0.02^2 = 0.0004$. Equation (22.7) gives

$$\sigma_n^2 = 0.9 \times 0.0001 + 0.1 \times 0.0004 = 0.00013$$

The estimate of the volatility, σ_n , for day n is therefore $\sqrt{0.00013}$, or 1.14%, per day. Note that the expected value of u_{n-1}^2 is σ_{n-1}^2 , or 0.0001. In this example, the realized value of u_{n-1}^2 is greater than the expected value, and as a result our volatility estimate increases. If the realized value of u_{n-1}^2 had been less than its expected value, our estimate of the volatility would have decreased.

The EWMA approach has the attractive feature that relatively little data need be stored. At any given time, only the current estimate of the variance rate and the most recent observation on the value of the market variable need be remembered. When a new observation on the market variable is obtained, a new daily percentage change is calculated and equation (22.7) is used to update the estimate of the variance rate. The old estimate of the variance rate and the old value of the market variable can then be discarded.

The EWMA approach is designed to track changes in the volatility. Suppose there is a big move in the market variable on day $n - 1$, so that u_{n-1}^2 is large. From equation (22.7) this causes the estimate of the current volatility to move upward. The value of λ governs how responsive the estimate of the daily volatility is to the most recent daily percentage change. A low value of λ leads to a great deal of weight being given to the u_{n-1}^2 when σ_n is calculated. In this case, the estimates produced for the volatility on successive days are themselves highly volatile. A high value of λ (i.e., a value close to 1.0) produces estimates of the daily volatility that respond relatively slowly to new information provided by the daily percentage change.

The RiskMetrics database, which was originally created by J.P. Morgan and made publicly available in 1994, uses the EWMA model with $\lambda = 0.94$ for updating daily volatility estimates in its RiskMetrics database. The company found that, across a range of different market variables, this value of λ gives forecasts of the variance rate that come closest to the realized variance rate.⁶ The realized variance rate on a particular day was calculated as an equally weighted average of the u_i^2 on the subsequent 25 days (see Problem 22.19).

⁶ See J.P. Morgan, *RiskMetrics Monitor*, Fourth Quarter, 1995. We will explain an alternative (maximum likelihood) approach to estimating parameters later in the chapter.

22.3 THE GARCH(1,1) MODEL

We now move on to discuss what is known as the GARCH(1,1) model, proposed by Bollerslev in 1986.⁷ The difference between the GARCH(1,1) model and the EWMA model is analogous to the difference between equation (22.4) and equation (22.5). In GARCH(1,1), σ_n^2 is calculated from a long-run average variance rate, V_L , as well as from σ_{n-1} and u_{n-1} . The equation for GARCH(1,1) is

$$\sigma_n^2 = \gamma V_L + \alpha u_{n-1}^2 + \beta \sigma_{n-1}^2 \quad (22.8)$$

where γ is the weight assigned to V_L , α is the weight assigned to u_{n-1}^2 , and β is the weight assigned to σ_{n-1}^2 . Since the weights must sum to unity, it follows that

$$\gamma + \alpha + \beta = 1$$

The EWMA model is a particular case of GARCH(1,1) where $\gamma = 0$, $\alpha = 1 - \lambda$, and $\beta = \lambda$.

The “(1,1)” in GARCH(1,1) indicates that σ_n^2 is based on the most recent observation of u^2 and the most recent estimate of the variance rate. The more general GARCH(p,q) model calculates σ_n^2 from the most recent p observations on u^2 and the most recent q estimates of the variance rate.⁸ GARCH(1,1) is by far the most popular of the GARCH models.

Setting $\omega = \gamma V_L$, the GARCH(1,1) model can also be written

$$\sigma_n^2 = \omega + \alpha u_{n-1}^2 + \beta \sigma_{n-1}^2 \quad (22.9)$$

This is the form of the model that is usually used for the purposes of estimating the parameters. Once ω , α , and β have been estimated, we can calculate γ as $1 - \alpha - \beta$. The long-term variance V_L can then be calculated as ω/γ . For a stable GARCH(1,1) process we require $\alpha + \beta < 1$. Otherwise the weight applied to the long-term variance is negative.

Example 22.2

Suppose that a GARCH(1,1) model is estimated from daily data as

$$\sigma_n^2 = 0.000002 + 0.13u_{n-1}^2 + 0.86\sigma_{n-1}^2$$

This corresponds to $\alpha = 0.13$, $\beta = 0.86$, and $\omega = 0.000002$. Because $\gamma = 1 - \alpha - \beta$, it follows that $\gamma = 0.01$. Because $\omega = \gamma V_L$, it follows that $V_L = 0.0002$. In other words, the long-run average variance per day implied by the model is 0.0002. This corresponds to a volatility of $\sqrt{0.0002} = 0.014$, or 1.4%, per day.

⁷ See T. Bollerslev, “Generalized Autoregressive Conditional Heteroscedasticity,” *Journal of Econometrics*, 31 (1986): 307–27.

⁸ Other GARCH models have been proposed that incorporate asymmetric news. These models are designed so that σ_n depends on the sign of u_{n-1} . Arguably, the models are more appropriate for equities than GARCH(1,1). As mentioned in Chapter 19, the volatility of an equity’s price tends to be inversely related to the price so that a negative u_{n-1} should have a bigger effect on σ_n than the same positive u_{n-1} . For a discussion of models for handling asymmetric news, see D. Nelson, “Conditional Heteroscedasticity and Asset Returns: A New Approach,” *Econometrica*, 59 (1990): 347–70; R. F. Engle and V. Ng, “Measuring and Testing the Impact of News on Volatility,” *Journal of Finance*, 48 (1993): 1749–78.

Suppose that the estimate of the volatility on day $n - 1$ is 1.6% per day, so that $\sigma_{n-1}^2 = 0.016^2 = 0.000256$, and that on day $n - 1$ the market variable decreased by 1%, so that $u_{n-1}^2 = 0.01^2 = 0.0001$. Then

$$\sigma_n^2 = 0.000002 + 0.13 \times 0.0001 + 0.86 \times 0.000256 = 0.00023516$$

The new estimate of the volatility is therefore $\sqrt{0.00023516} = 0.0153$, or 1.53%, per day.

The Weights

Substituting for σ_{n-1}^2 in equation (22.9) gives

$$\sigma_n^2 = \omega + \alpha u_{n-1}^2 + \beta(\omega + \alpha u_{n-2}^2 + \beta \sigma_{n-2}^2)$$

or

$$\sigma_n^2 = \omega + \beta\omega + \alpha u_{n-1}^2 + \alpha\beta u_{n-2}^2 + \beta^2 \sigma_{n-2}^2$$

Substituting for σ_{n-2}^2 gives

$$\sigma_n^2 = \omega + \beta\omega + \beta^2\omega + \alpha u_{n-1}^2 + \alpha\beta u_{n-2}^2 + \alpha\beta^2 u_{n-3}^2 + \beta^3 \sigma_{n-3}^2$$

Continuing in this way, we see that the weight applied to u_{n-i}^2 is $\alpha\beta^{i-1}$. The weights decline exponentially at rate β . The parameter β can be interpreted as a “decay rate”. It is similar to λ in the EWMA model. It defines the relative importance of the observations on the u 's in determining the current variance rate. For example, if $\beta = 0.9$, then u_{n-2}^2 is only 90% as important as u_{n-1}^2 ; u_{n-3}^2 is 81% as important as u_{n-1}^2 ; and so on. The GARCH(1,1) model is similar to the EWMA model except that, in addition to assigning weights that decline exponentially to past u^2 , it also assigns some weight to the long-run average volatility.

Mean Reversion

The GARCH (1,1) model recognizes that over time the variance tends to get pulled back to a long-run average level of V_L . The amount of weight assigned to V_L is $\gamma = 1 - \alpha - \beta$. The GARCH(1,1) is equivalent to a model where the variance V follows the stochastic process

$$dV = a(V_L - V) dt + \xi V dz$$

where time is measured in days, $a = 1 - \alpha - \beta$, and $\xi = \alpha\sqrt{2}$ (see Problem 22.14). This is a mean-reverting model. The variance has a drift that pulls it back to V_L at rate a . When $V > V_L$, the variance has a negative drift; when $V < V_L$, it has a positive drift. Superimposed on the drift is a volatility ξ . Chapter 26 discusses this type of model further.

22.4 CHOOSING BETWEEN THE MODELS

In practice, variance rates do tend to be mean reverting. The GARCH(1,1) model incorporates mean reversion, whereas the EWMA model does not. GARCH (1,1) is therefore theoretically more appealing than the EWMA model.

In the next section, we will discuss how best-fit parameters ω , α , and β in GARCH(1, 1) can be estimated. When the parameter ω is zero, the GARCH(1, 1) reduces to EWMA. In circumstances where the best-fit value of ω turns out to be negative, the GARCH(1, 1) model is not stable and it makes sense to switch to the EWMA model.

22.5 MAXIMUM LIKELIHOOD METHODS

It is now appropriate to discuss how the parameters in the models we have been considering are estimated from historical data. The approach used is known as the *maximum likelihood method*. It involves choosing values for the parameters that maximize the chance (or likelihood) of the data occurring.

To illustrate the method, we start with a very simple example. Suppose that we sample 10 stocks at random on a certain day and find that the price of one of them declined on that day and the prices of the other nine either remained the same or increased. What is the best estimate of the probability of a price decline? The natural answer is 0.1. Let us see if this is what the maximum likelihood method gives.

Suppose that the probability of a price decline is p . The probability that one particular stock declines in price and the other nine do not is $p(1-p)^9$. Using the maximum likelihood approach, the best estimate of p is the one that maximizes $p(1-p)^9$. Differentiating this expression with respect to p and setting the result equal to zero, we find that $p = 0.1$ maximizes the expression. This shows that the maximum likelihood estimate of p is 0.1, as expected.

Estimating a Constant Variance

Our next example of maximum likelihood methods considers the problem of estimating the variance of a variable X from m observations on X when the underlying distribution is normal with zero mean. Assume that the observations are u_1, u_2, \dots, u_m . Denote the variance by v . The likelihood of u_i being observed is defined as the probability density function for X when $X = u_i$. This is

$$\frac{1}{\sqrt{2\pi v}} \exp\left(\frac{-u_i^2}{2v}\right)$$

The likelihood of m observations occurring in the order in which they are observed is

$$\prod_{i=1}^m \left[\frac{1}{\sqrt{2\pi v}} \exp\left(\frac{-u_i^2}{2v}\right) \right] \quad (22.10)$$

Using the maximum likelihood method, the best estimate of v is the value that maximizes this expression.

Maximizing an expression is equivalent to maximizing the logarithm of the expression. Taking logarithms of the expression in equation (22.10) and ignoring constant multiplicative factors, it can be seen that we wish to maximize

$$\sum_{i=1}^m \left[-\ln(v) - \frac{u_i^2}{v} \right] \quad (22.11)$$

or

$$-m \ln(v) - \sum_{i=1}^m \frac{u_i^2}{v}$$

Differentiating this expression with respect to v and setting the resulting equation to zero, we see that the maximum likelihood estimator of v is⁹

$$\frac{1}{m} \sum_{i=1}^m u_i^2$$

Estimating GARCH (1,1) Parameters

We now consider how the maximum likelihood method can be used to estimate the parameters when GARCH (1,1) or some other volatility updating scheme is used. Define $v_i = \sigma_i^2$ as the variance estimated for day i . Assume that the probability distribution of u_i conditional on the variance is normal. A similar analysis to the one just given shows the best parameters are the ones that maximize

$$\prod_{i=1}^m \left[\frac{1}{\sqrt{2\pi v_i}} \exp\left(\frac{-u_i^2}{2v_i}\right) \right]$$

Taking logarithms, we see that this is equivalent to maximizing

$$\sum_{i=1}^m \left[-\ln(v_i) - \frac{u_i^2}{v_i} \right] \quad (22.12)$$

This is the same as the expression in equation (22.11), except that v is replaced by v_i . It is necessary to search iteratively to find the parameters in the model that maximize the expression in equation (22.12).

The spreadsheet in Table 22.1 indicates how the calculations could be organized for the GARCH(1,1) model. The table analyzes data on the S&P 500 between July 18, 2005, and August 13, 2010.¹⁰ The numbers in the table are based on trial estimates of the three GARCH(1,1) parameters: ω , α , and β . The first column in the table records the date. The second column counts the days. The third column shows the S&P 500, S_i , at the end of day i . The fourth column shows the proportional change in the S&P 500 between the end of day $i - 1$ and the end of day i . This is $u_i = (S_i - S_{i-1})/S_{i-1}$. The fifth column shows the estimate of the variance rate, $v_i = \sigma_i^2$, for day i made at the end of day $i - 1$. On day 3, we start things off by setting the variance equal to u_2^2 . On subsequent days, equation (22.9) is used. The sixth column tabulates the likelihood measure, $-\ln(v_i) - u_i^2/v_i$. The values in the fifth and sixth columns are based on the current trial estimates of ω , α , and β . We are interested in choosing ω , α , and β to maximize the sum of the numbers in the sixth column. This involves an iterative search procedure.¹¹

⁹ This confirms the point made in footnote 3.

¹⁰ The data and calculations can be found at www.rotman.utoronto.ca/~hull/OFOD/GarchExample.

¹¹ As discussed later, a general purpose algorithm such as Solver in Microsoft's Excel can be used. Alternatively, a special purpose algorithm, such as Levenberg–Marquardt, can be used. See, e.g., W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical Recipes in C: The Art of Scientific Computing*, Cambridge University Press, 1988.

Table 22.1 Estimation of Parameters in GARCH(1,1) Model for S&P 500 between July 18, 2005, and August 13, 2010.

Date	Day i	S_i	u_i	$v_i = \sigma_i^2$	$-\ln(v_i) - u_i^2/v_i$
18-Jul-2005	1	1221.13			
19-Jul-2005	2	1229.35	0.006731		
20-Jul-2005	3	1235.20	0.004759	0.00004531	9.5022
21-Jul-2005	4	1227.04	-0.006606	0.00004447	9.0393
22-Jul-2005	5	1233.68	0.005411	0.00004546	9.3545
25-Jul-2005	6	1229.03	-0.003769	0.00004517	9.6906
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
11-Aug-2010	1277	1089.47	-0.028179	0.00011834	2.3322
12-Aug-2010	1278	1083.61	-0.005379	0.00017527	8.4841
13-Aug-2010	1279	1079.25	-0.004024	0.00016327	8.6209
					10,228.2349

Trial estimates of GARCH parameters

$$\omega = 0.000001347 \quad \alpha = 0.08339 \quad \beta = 0.9101$$

In our example, the optimal values of the parameters turn out to be

$$\omega = 0.000001366, \quad \alpha = 0.083394, \quad \beta = 0.910116$$

and the maximum value of the function in equation (22.12) is 10,228.2349. The numbers shown in Table 22.1 were calculated on the final iteration of the search for the optimal ω , α , and β .

The long-term variance rate, V_L , in our example is

$$\frac{\omega}{1 - \alpha - \beta} = \frac{0.000001366}{0.006490} = 0.0002075$$

The long-term volatility is $\sqrt{0.0002075}$, or 1.4404%, per day.

Figures 22.1 and 22.2 show the S&P 500 index and its GARCH(1,1) volatility during the 5-year period covered by the data. Most of the time, the volatility was less than 2% per day, but volatilities as high as 5% per day were experienced during the credit crisis. (The very high volatilities are also indicated by the VIX index—see Section 14.11.)

An alternative approach to estimating parameters in GARCH(1,1), which is sometimes more robust, is known as *variance targeting*.¹² This involves setting the long-run average variance rate, V_L , equal to the sample variance calculated from the data (or to some other value that is believed to be reasonable). The value of ω then equals $V_L(1 - \alpha - \beta)$ and only two parameters have to be estimated. For the data in Table 22.1, the sample variance is 0.0002412, which gives a daily volatility of 1.5531%. Setting V_L equal to the sample variance, the values of α and β that maximize the objective function in equation (22.12) are 0.08445 and 0.9101, respectively. The value of the objective function is 10,228.1941, only marginally below the value of 10,228.2349 obtained using the earlier procedure.

¹² See R. Engle and J. Mezrich, “GARCH for Groups,” *Risk*, August 1996: 36–40.

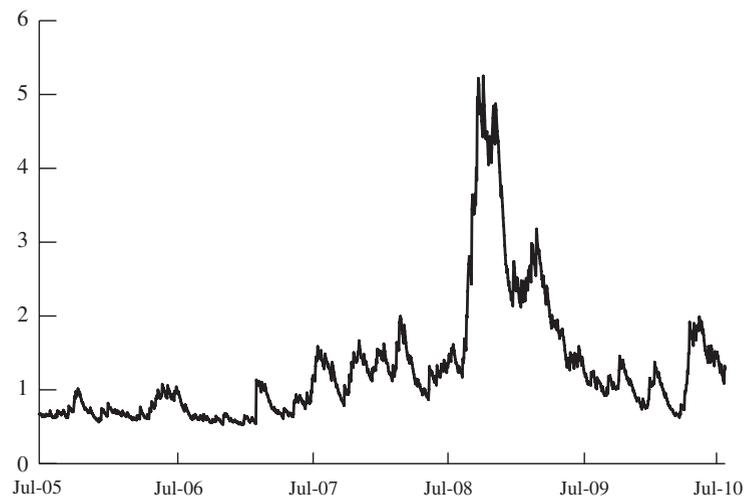
Figure 22.1 S&P 500 index: July 18, 2005, to August 13, 2010.



When the EWMA model is used, the estimation procedure is relatively simple. We set $\omega = 0$, $\alpha = 1 - \lambda$, and $\beta = \lambda$, and only one parameter has to be estimated. In the data in Table 22.1, the value of λ that maximizes the objective function in equation (22.12) is 0.9374 and the value of the objective function is 10,192.5104.

Both GARCH (1,1) and the EWMA method can be implemented by using the Solver routine in Excel to search for the values of the parameters that maximize the likelihood function. The routine works well provided that the spreadsheet is structured so that the parameters being searched for have roughly equal values. For example, in GARCH (1,1)

Figure 22.2 Daily volatility of S&P 500 index: July 18, 2005, to August 13, 2010.



we could let cells A1, A2, and A3 contain $\omega \times 10^5$, 10α , and β . We could then set $B1 = A1/100,000$, $B2 = A2/10$, and $B3 = A3$. We would use B1, B2, and B3 to calculate the likelihood function. We would ask Solver to calculate the values of A1, A2, and A3 that maximize the likelihood function. Occasionally Solver gives a local maximum, so testing a number of different starting values for parameters is a good idea.

How Good Is the Model?

The assumption underlying a GARCH model is that volatility changes with the passage of time. During some periods volatility is relatively high; during other periods it is relatively low. To put this another way, when u_i^2 is high, there is a tendency for u_{i+1}^2 , u_{i+2}^2, \dots to be high; when u_i^2 is low, there is a tendency for $u_{i+1}^2, u_{i+2}^2, \dots$ to be low. We can test how true this is by examining the autocorrelation structure of the u_i^2 .

Let us assume the u_i^2 do exhibit autocorrelation. If a GARCH model is working well, it should remove the autocorrelation. We can test whether it has done so by considering the autocorrelation structure for the variables u_i^2/σ_i^2 . If these show very little autocorrelation, our model for σ_i has succeeded in explaining autocorrelations in the u_i^2 .

Table 22.2 shows results for the S&P 500 data used above. The first column shows the lags considered when the autocorrelation is calculated. The second shows autocorrelations for u_i^2 ; the third shows autocorrelations for u_i^2/σ_i^2 .¹³ The table shows that the autocorrelations are positive for u_i^2 for all lags between 1 and 15. In the case of u_i^2/σ_i^2 , some of the autocorrelations are positive and some are negative. They are all much smaller in magnitude than the autocorrelations for u_i^2 .

Table 22.2 Autocorrelations before and after the use of a GARCH model for S&P 500 data.

<i>Time lag</i>	<i>Autocorrelation for u_i^2</i>	<i>Autocorrelation for u_i^2/σ_i^2</i>
1	0.183	-0.063
2	0.385	-0.004
3	0.160	-0.007
4	0.301	0.022
5	0.339	0.014
6	0.308	-0.011
7	0.329	0.026
8	0.207	0.038
9	0.324	0.041
10	0.269	0.083
11	0.431	-0.007
12	0.286	0.006
13	0.224	0.001
14	0.121	0.017
15	0.222	-0.031

¹³ For a series x_i , the autocorrelation with a lag of k is the coefficient of correlation between x_i and x_{i+k} .

The GARCH model appears to have done a good job in explaining the data. For a more scientific test, we can use what is known as the Ljung–Box statistic.¹⁴ If a certain series has m observations the Ljung–Box statistic is

$$m \sum_{k=1}^K w_k \eta_k^2$$

where η_k is the autocorrelation for a lag of k , K is the number of lags considered, and

$$w_k = \frac{m+2}{m-k}$$

For $K = 15$, zero autocorrelation can be rejected with 95% confidence when the Ljung–Box statistic is greater than 25.

From Table 22.2, the Ljung–Box statistic for the u_i^2 series is about 1,566. This is strong evidence of autocorrelation. For the u_i^2/σ_i^2 series, the Ljung–Box statistic is 21.7, suggesting that the autocorrelation has been largely removed by the GARCH model.

22.6 USING GARCH(1,1) TO FORECAST FUTURE VOLATILITY

The variance rate estimated at the end of day $n - 1$ for day n , when GARCH(1,1) is used, is

$$\sigma_n^2 = (1 - \alpha - \beta)V_L + \alpha u_{n-1}^2 + \beta \sigma_{n-1}^2$$

so that

$$\sigma_n^2 - V_L = \alpha(u_{n-1}^2 - V_L) + \beta(\sigma_{n-1}^2 - V_L)$$

On day $n + t$ in the future,

$$\sigma_{n+t}^2 - V_L = \alpha(u_{n+t-1}^2 - V_L) + \beta(\sigma_{n+t-1}^2 - V_L)$$

The expected value of u_{n+t-1}^2 is σ_{n+t-1}^2 . Hence,

$$E[\sigma_{n+t}^2 - V_L] = (\alpha + \beta)E[\sigma_{n+t-1}^2 - V_L]$$

where E denotes expected value. Using this equation repeatedly yields

$$E[\sigma_{n+t}^2 - V_L] = (\alpha + \beta)^t (\sigma_n^2 - V_L)$$

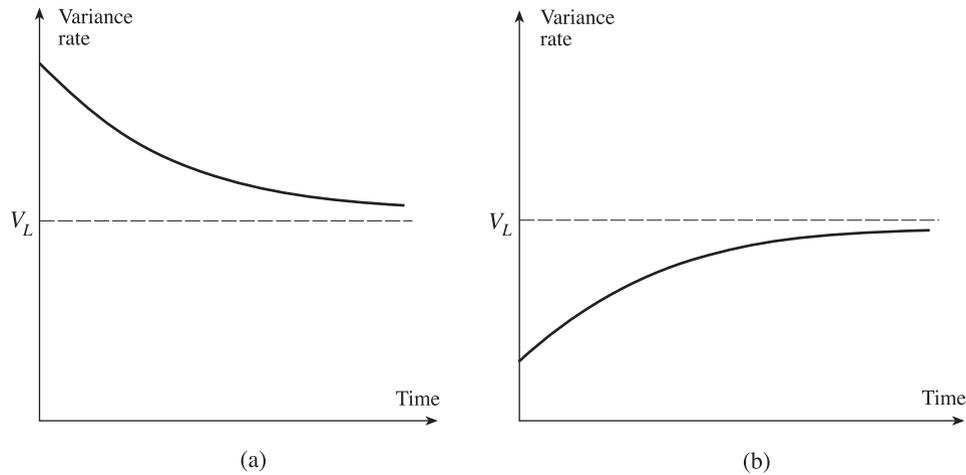
or

$$E[\sigma_{n+t}^2] = V_L + (\alpha + \beta)^t (\sigma_n^2 - V_L) \quad (22.13)$$

This equation forecasts the volatility on day $n + t$ using the information available at the end of day $n - 1$. In the EWMA model, $\alpha + \beta = 1$ and equation (22.13) shows that the expected future variance rate equals the current variance rate. When $\alpha + \beta < 1$, the final term in the equation becomes progressively smaller as t increases. Figure 22.3 shows the expected path followed by the variance rate for situations where the current variance rate is different from V_L . As mentioned earlier, the variance rate exhibits mean reversion with a reversion level of V_L and a reversion rate of $1 - \alpha - \beta$. Our forecast of the future

¹⁴ See G. M. Ljung and G. E. P. Box, “On a Measure of Lack of Fit in Time Series Models,” *Biometrika*, 65 (1978): 297–303.

Figure 22.3 Expected path for the variance rate when (a) current variance rate is above long-term variance rate and (b) current variance rate is below long-term variance rate.



variance rate tends towards V_L as we look further and further ahead. This analysis emphasizes the point that we must have $\alpha + \beta < 1$ for a stable GARCH(1,1) process. When $\alpha + \beta > 1$, the weight given to the long-term average variance is negative and the process is “mean fleeing” rather than “mean reverting”.

For the S&P 500 data considered earlier, $\alpha + \beta = 0.9935$ and $V_L = 0.0002075$. Suppose that the estimate of the current variance rate per day is 0.0003. (This corresponds to a volatility of 1.732% per day.) In 10 days, the expected variance rate is

$$0.0002075 + 0.9935^{10}(0.0003 - 0.0002075) = 0.0002942$$

The expected volatility per day is 1.72%, still well above the long-term volatility of 1.44% per day. However, the expected variance rate in 500 days is

$$0.0002075 + 0.9935^{500}(0.0003 - 0.0002075) = 0.0002110$$

and the expected volatility per day is 1.45%, very close to the long-term volatility.

Volatility Term Structures

Suppose it is day n . Define:

$$V(t) = E(\sigma_{n+t}^2)$$

and

$$a = \ln \frac{1}{\alpha + \beta}$$

so that equation (22.13) becomes

$$V(t) = V_L + e^{-at}[V(0) - V_L]$$

Here, $V(t)$ is an estimate of the instantaneous variance rate in t days. The average

Table 22.3 S&P 500 volatility term structure predicted from GARCH(1, 1).

Option life (days)	10	30	50	100	500
Option volatility (% per annum)	27.36	27.10	26.87	26.35	24.32

variance rate per day between today and time T is given by

$$\frac{1}{T} \int_0^T V(t) dt = V_L + \frac{1 - e^{-aT}}{aT} [V(0) - V_L]$$

The larger T is, the closer this is to V_L . Define $\sigma(T)$ as the volatility per annum that should be used to price a T -day option under GARCH(1, 1). Assuming 252 days per year, $\sigma(T)^2$ is 252 times the average variance rate per day, so that

$$\sigma(T)^2 = 252 \left(V_L + \frac{1 - e^{-aT}}{aT} [V(0) - V_L] \right) \quad (22.14)$$

As discussed in Chapter 19, the market prices of different options on the same asset are often used to calculate a *volatility term structure*. This is the relationship between the implied volatilities of the options and their maturities. Equation (22.14) can be used to estimate a volatility term structure based on the GARCH(1, 1) model. The estimated volatility term structure is not usually the same as the actual volatility term structure. However, as we will show, it is often used to predict the way that the actual volatility term structure will respond to volatility changes.

When the current volatility is above the long-term volatility, the GARCH(1, 1) model estimates a downward-sloping volatility term structure. When the current volatility is below the long-term volatility, it estimates an upward-sloping volatility term structure. In the case of the S&P 500 data, $a = \ln(1/0.99351) = 0.006511$ and $V_L = 0.0002075$. Suppose that the current variance rate per day, $V(0)$, is estimated as 0.0003 per day. It follows from equation (22.14) that

$$\sigma(T)^2 = 252 \left(0.0002075 + \frac{1 - e^{-0.006511T}}{0.006511T} (0.0003 - 0.0002075) \right)$$

where T is measured in days. Table 22.3 shows the volatility per year for different values of T .

Impact of Volatility Changes

Equation (22.14) can be written

$$\sigma(T)^2 = 252 \left[V_L + \frac{1 - e^{-aT}}{aT} \left(\frac{\sigma(0)^2}{252} - V_L \right) \right]$$

When $\sigma(0)$ changes by $\Delta\sigma(0)$, $\sigma(T)$ changes by approximately

$$\frac{1 - e^{-aT}}{aT} \frac{\sigma(0)}{\sigma(T)} \Delta\sigma(0) \quad (22.15)$$

Table 22.4 Impact of 1% change in the instantaneous volatility predicted from GARCH(1,1).

Option life (days)	10	30	50	100	500
Increase in volatility (%)	0.97	0.92	0.87	0.77	0.33

Table 22.4 shows the effect of a volatility change on options of varying maturities for the S&P 500 data considered above. We assume as before that $V(0) = 0.0003$, so that $\sigma(0) = \sqrt{252} \times \sqrt{0.0003} = 27.50\%$. The table considers a 100-basis-point change in the instantaneous volatility from 27.50% per year to 28.50% per year. This means that $\Delta\sigma(0) = 0.01$, or 1%.

Many financial institutions use analyses such as this when determining the exposure of their books to volatility changes. Rather than consider an across-the-board increase of 1% in implied volatilities when calculating vega, they relate the size of the volatility increase that is considered to the maturity of the option. Based on Table 22.4, a 0.97% volatility increase would be considered for a 10-day option, a 0.92% increase for a 30-day option, a 0.87% increase for a 50-day option, and so on.

22.7 CORRELATIONS

The discussion so far has centered on the estimation and forecasting of volatility. As explained in Chapter 21, correlations also play a key role in the calculation of VaR. In this section, we show how correlation estimates can be updated in a similar way to volatility estimates.

The correlation between two variables X and Y can be defined as

$$\frac{\text{cov}(X, Y)}{\sigma_X \sigma_Y}$$

where σ_X and σ_Y are the standard deviations of X and Y and $\text{cov}(X, Y)$ is the covariance between X and Y . The covariance between X and Y is defined as

$$E[(X - \mu_X)(Y - \mu_Y)]$$

where μ_X and μ_Y are the means of X and Y , and E denotes the expected value. Although it is easier to develop intuition about the meaning of a correlation than it is for a covariance, it is covariances that are the fundamental variables of our analysis.¹⁵

Define x_i and y_i as the percentage changes in X and Y between the end of day $i - 1$ and the end of day i :

$$x_i = \frac{X_i - X_{i-1}}{X_{i-1}}, \quad y_i = \frac{Y_i - Y_{i-1}}{Y_{i-1}}$$

where X_i and Y_i are the values of X and Y at the end of day i . We also define the

¹⁵ An analogy here is that variance rates were the fundamental variables for the EWMA and GARCH procedures in the first part of this chapter, even though volatilities are easier to understand.

following:

$\sigma_{x,n}$: Daily volatility of variable X , estimated for day n

$\sigma_{y,n}$: Daily volatility of variable Y , estimated for day n

cov_n : Estimate of covariance between daily changes in X and Y , calculated on day n .

The estimate of the correlation between X and Y on day n is

$$\frac{\text{cov}_n}{\sigma_{x,n} \sigma_{y,n}}$$

Using equal weighting and assuming that the means of x_i and y_i are zero, equation (22.3) shows that the variance rates of X and Y can be estimated from the most recent m observations as

$$\sigma_{x,n}^2 = \frac{1}{m} \sum_{i=1}^m x_{n-i}^2, \quad \sigma_{y,n}^2 = \frac{1}{m} \sum_{i=1}^m y_{n-i}^2$$

A similar estimate for the covariance between X and Y is

$$\text{cov}_n = \frac{1}{m} \sum_{i=1}^m x_{n-i} y_{n-i} \tag{22.16}$$

One alternative for updating covariances is an EWMA model similar to equation (22.7). The formula for updating the covariance estimate is then

$$\text{cov}_n = \lambda \text{cov}_{n-1} + (1 - \lambda)x_{n-1} y_{n-1}$$

A similar analysis to that presented for the EWMA volatility model shows that the weights given to observations on the $x_i y_i$ decline as we move back through time. The lower the value of λ , the greater the weight that is given to recent observations.

Example 22.3

Suppose that $\lambda = 0.95$ and that the estimate of the correlation between two variables X and Y on day $n - 1$ is 0.6. Suppose further that the estimate of the volatilities for the X and Y on day $n - 1$ are 1% and 2%, respectively. From the relationship between correlation and covariance, the estimate of the covariance between the X and Y on day $n - 1$ is

$$0.6 \times 0.01 \times 0.02 = 0.00012$$

Suppose that the percentage changes in X and Y on day $n - 1$ are 0.5% and 2.5%, respectively. The variance and covariance for day n would be updated as follows:

$$\sigma_{x,n}^2 = 0.95 \times 0.01^2 + 0.05 \times 0.005^2 = 0.00009625$$

$$\sigma_{y,n}^2 = 0.95 \times 0.02^2 + 0.05 \times 0.025^2 = 0.00041125$$

$$\text{cov}_n = 0.95 \times 0.00012 + 0.05 \times 0.005 \times 0.025 = 0.00012025$$

The new volatility of X is $\sqrt{0.00009625} = 0.981\%$ and the new volatility of Y is $\sqrt{0.00041125} = 2.028\%$. The new coefficient of correlation between X and Y is

$$\frac{0.00012025}{0.00981 \times 0.02028} = 0.6044$$

GARCH models can also be used for updating covariance estimates and forecasting the future level of covariances. For example, the GARCH(1,1) model for updating a covariance is

$$\text{cov}_n = \omega + \alpha x_{n-1} y_{n-1} + \beta \text{cov}_{n-1}$$

and the long-term average covariance is $\omega/(1 - \alpha - \beta)$. Formulas similar to those in equations (22.13) and (22.14) can be developed for forecasting future covariances and calculating the average covariance during the life of an option.¹⁶

Consistency Condition for Covariances

Once all the variances and covariances have been calculated, a variance–covariance matrix can be constructed. As explained in Section 21.4, when $i \neq j$, the (i, j) th element of this matrix shows the covariance between variable i and variable j . When $i = j$, it shows the variance of variable i .

Not all variance–covariance matrices are internally consistent. The condition for an $N \times N$ variance–covariance matrix Ω to be internally consistent is

$$\mathbf{w} \Omega \mathbf{w} \geq 0 \tag{22.17}$$

for all $N \times 1$ vectors \mathbf{w} , where \mathbf{w} is the transpose of \mathbf{w} . A matrix that satisfies this property is known as *positive-semidefinite*.

To understand why the condition in equation (22.17) must hold, suppose that \mathbf{w} is $[w_1, w_2, \dots, w_n]$. The expression $\mathbf{w} \Omega \mathbf{w}$ is the variance of $w_1 x_1 + w_2 x_2 + \dots + w_n x_n$, where x_i is the value of variable i . As such, it cannot be negative.

To ensure that a positive-semidefinite matrix is produced, variances and covariances should be calculated consistently. For example, if variances are calculated by giving equal weight to the last m data items, the same should be done for covariances. If variances are updated using an EWMA model with $\lambda = 0.94$, the same should be done for covariances.

An example of a variance–covariance matrix that is not internally consistent is

$$\begin{bmatrix} 1 & 0 & 0.9 \\ 0 & 1 & 0.9 \\ 0.9 & 0.9 & 1 \end{bmatrix}$$

The variance of each variable is 1.0, and so the covariances are also coefficients of correlation. The first variable is highly correlated with the third variable and the second variable is highly correlated with the third variable. However, there is no correlation at all between the first and second variables. This seems strange. When \mathbf{w} is set equal to $(1, 1, -1)$, the condition in equation (22.17) is not satisfied, proving that the matrix is not positive-semidefinite.¹⁷

¹⁶ The ideas in this chapter can be extended to multivariate GARCH models, where an entire variance–covariance matrix is updated in a consistent way. For a discussion of alternative approaches, see R. Engle and J. Mezrich, “GARCH for Groups,” *Risk*, August 1996: 36–40.

¹⁷ It can be shown that the condition for a 3×3 matrix of correlations to be internally consistent is

$$\rho_{12}^2 + \rho_{13}^2 + \rho_{23}^2 - 2\rho_{12}\rho_{13}\rho_{23} \leq 1$$

where ρ_{ij} is the coefficient of correlation between variables i and j .

22.8 APPLICATION OF EWMA TO FOUR-INDEX EXAMPLE

We now return to the example considered in Section 21.2. This involved a portfolio on September 25, 2008, consisting of a \$4 million investment in the Dow Jones Industrial Average, a \$3 million investment in the FTSE 100, a \$1 million investment in the CAC 40, and a \$2 million investment in the Nikkei 225. Daily returns were collected over 500 days ending on September 25, 2008. Data and all calculations presented here can be found at: www.rotman.utoronto.ca/~hull/OFOD/VaRExample.

The correlation matrix that would be calculated on September 25, 2008, by giving equal weight to the last 500 returns is shown in Table 22.5. The FTSE 100 and CAC 40 are very highly correlated. The Dow Jones Industrial Average is moderately highly correlated with both the FTSE 100 and the CAC 40. The correlation of the Nikkei 225 with other indices is less high.

The covariance matrix for the equal-weight case is shown in Table 22.6. From equation (21.3), this matrix gives the variance of the portfolio losses (\$000s) as 8,761.833. The standard deviation is the square root of this, or 93.60. The one-day 99% VaR in \$000s is therefore $2.33 \times 93.60 = 217.757$. This is \$217,757, which compares with \$253,385, calculated using the historical simulation approach in Section 21.2.

Instead of calculating variances and covariances by giving equal weight to all observed returns, we now use the exponentially weighted moving average method with $\lambda = 0.94$. This gives the variance–covariance matrix in Table 22.7.¹⁸ From equation (21.3), the

Table 22.5 Correlation matrix on September 25, 2008, calculated by giving equal weight to the last 500 daily returns: variable 1 is DJIA; variable 2 is FTSE 100; variable 3 is CAC 40; variable 4 is Nikkei 225.

$$\begin{bmatrix} 1 & 0.489 & 0.496 & -0.062 \\ 0.489 & 1 & 0.918 & 0.201 \\ 0.496 & 0.918 & 1 & 0.211 \\ -0.062 & 0.201 & 0.211 & 1 \end{bmatrix}$$

Table 22.6 Covariance matrix on September 25, 2008, calculated by giving equal weight to the last 500 daily returns: variable 1 is DJIA; variable 2 is FTSE 100; variable 3 is CAC 40; variable 4 is Nikkei 225.

$$\begin{bmatrix} 0.0001227 & 0.0000768 & 0.0000767 & -0.0000095 \\ 0.0000768 & 0.0002010 & 0.0001817 & 0.0000394 \\ 0.0000767 & 0.0001817 & 0.0001950 & 0.0000407 \\ -0.0000095 & 0.0000394 & 0.0000407 & 0.0001909 \end{bmatrix}$$

¹⁸ In the EWMA calculations, the variance is initially set equal to the population variance. But all reasonable starting variances give essentially the same result because in this case all we are interested in is the final variance.

Table 22.7 Covariance matrix on September 25, 2008, calculated using the EWMA method with $\lambda = 0.94$: variable 1 is DJIA; variable 2 is FTSE 100; variable 3 is CAC 40; variable 4 is Nikkei 225.

$$\begin{bmatrix} 0.0004801 & 0.0004303 & 0.0004257 & -0.0000396 \\ 0.0004303 & 0.0010314 & 0.0009630 & 0.0002095 \\ 0.0004257 & 0.0009630 & 0.0009535 & 0.0001681 \\ -0.0000396 & 0.0002095 & 0.0001681 & 0.0002541 \end{bmatrix}$$

variance of portfolio losses (\$000s) is 40,995.765. The standard deviation is the square root of this, or 202.474. The one-day 99% VaR is therefore

$$2.33 \times 202.474 = 471.025$$

This is \$471,025, over twice as high as the value given when returns are equally weighted. Tables 22.8 and 22.9 show the reasons. The standard deviation of a portfolio consisting of long positions in securities increases with the standard deviations of security returns and also with the correlations between security returns. Table 22.8 shows that the estimated daily standard deviations are much higher when EWMA is used than when data are equally weighted. This is because volatilities were much higher during the period immediately preceding September 25, 2008, than during the rest of the 500 days covered by the data. Comparing Table 22.9 with Table 22.5, we see that correlations had also increased.¹⁹

Table 22.8 Volatilities (% per day) using equal weighting and EWMA.

	<i>DJIA</i>	<i>FTSE 100</i>	<i>CAC 40</i>	<i>Nikkei 225</i>
Equal weighting:	1.11	1.42	1.40	1.38
EWMA:	2.19	3.21	3.09	1.59

Table 22.9 Correlation matrix on September 25, 2008, calculated using the EWMA method: variable 1 is DJIA; variable 2 is FTSE 100; variable 3 is CAC 40; variable 4 is Nikkei 225.

$$\begin{bmatrix} 1 & 0.611 & 0.629 & -0.113 \\ 0.611 & 1 & 0.971 & 0.409 \\ 0.629 & 0.971 & 1 & 0.342 \\ -0.113 & 0.409 & 0.342 & 1 \end{bmatrix}$$

¹⁹ This is an example of the phenomenon that correlations tend to increase in adverse market conditions.

SUMMARY

Most popular option pricing models, such as Black–Scholes, assume that the volatility of the underlying asset is constant. This assumption is far from perfect. In practice, the volatility of an asset, like the asset's price, is a stochastic variable. Unlike the asset price, it is not directly observable. This chapter has discussed procedures for attempting to keep track of the current level of volatility.

We define u_i as the percentage change in a market variable between the end of day $i - 1$ and the end of day i . The variance rate of the market variable (that is, the square of its volatility) is calculated as a weighted average of the u_i^2 . The key feature of the procedures that have been discussed here is that they do not give equal weight to the observations on the u_i^2 . The more recent an observation, the greater the weight assigned to it. In the EWMA and the GARCH(1,1) models, the weights assigned to observations decrease exponentially as the observations become older. The GARCH(1,1) model differs from the EWMA model in that some weight is also assigned to the long-run average variance rate. It has a structure that enables forecasts of the future level of variance rate to be produced relatively easily.

Maximum likelihood methods are usually used to estimate parameters from historical data in the EWMA, GARCH(1,1), and similar models. These methods involve using an iterative procedure to determine the parameter values that maximize the chance or likelihood that the historical data will occur. Once its parameters have been determined, a GARCH(1,1) model can be judged by how well it removes autocorrelation from the u_i^2 .

For every model that is developed to track variances, there is a corresponding model that can be developed to track covariances. The procedures described here can therefore be used to update the complete variance–covariance matrix used in value at risk calculations.

FURTHER READING

- Bollerslev, T. “Generalized Autoregressive Conditional Heteroscedasticity,” *Journal of Econometrics*, 31 (1986): 307–27.
- Cumby, R., S. Figlewski, and J. Hasbrook. “Forecasting Volatilities and Correlations with EGARCH Models,” *Journal of Derivatives*, 1, 2 (Winter 1993): 51–63.
- Engle, R. F. “Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of UK Inflation,” *Econometrica* 50 (1982): 987–1008.
- Engle R. F., and J. Mezrich. “Grappling with GARCH,” *Risk*, September 1995: 112–117.
- Engle, R. F., and J. Mezrich, “GARCH for Groups,” *Risk*, August 1996: 36–40.
- Engle, R. F., and V. Ng, “Measuring and Testing the Impact of News on Volatility,” *Journal of Finance*, 48 (1993): 1749–78.
- Nelson, D. “Conditional Heteroscedasticity and Asset Returns: A New Approach,” *Econometrica*, 59 (1990): 347–70.
- Noh, J., R. F. Engle, and A. Kane. “Forecasting Volatility and Option Prices of the S&P 500 Index,” *Journal of Derivatives*, 2 (1994): 17–30.

Practice Questions (Answers in Solutions Manual)

- 22.1. Explain the exponentially weighted moving average (EWMA) model for estimating volatility from historical data.
- 22.2. What is the difference between the exponentially weighted moving average model and the GARCH(1,1) model for updating volatilities?
- 22.3. The most recent estimate of the daily volatility of an asset is 1.5% and the price of the asset at the close of trading yesterday was \$30.00. The parameter λ in the EWMA model is 0.94. Suppose that the price of the asset at the close of trading today is \$30.50. How will this cause the volatility to be updated by the EWMA model?
- 22.4. A company uses an EWMA model for forecasting volatility. It decides to change the parameter λ from 0.95 to 0.85. Explain the likely impact on the forecasts.
- 22.5. The volatility of a certain market variable is 30% per annum. Calculate a 99% confidence interval for the size of the percentage daily change in the variable.
- 22.6. A company uses the GARCH(1,1) model for updating volatility. The three parameters are ω , α , and β . Describe the impact of making a small increase in each of the parameters while keeping the others fixed.
- 22.7. The most recent estimate of the daily volatility of the US dollar/sterling exchange rate is 0.6% and the exchange rate at 4 p.m. yesterday was 1.5000. The parameter λ in the EWMA model is 0.9. Suppose that the exchange rate at 4 p.m. today proves to be 1.4950. How would the estimate of the daily volatility be updated?
- 22.8. Assume that S&P 500 at close of trading yesterday was 1,040 and the daily volatility of the index was estimated as 1% per day at that time. The parameters in a GARCH(1,1) model are $\omega = 0.000002$, $\alpha = 0.06$, and $\beta = 0.92$. If the level of the index at close of trading today is 1,060, what is the new volatility estimate?
- 22.9. Suppose that the daily volatilities of asset A and asset B, calculated at the close of trading yesterday, are 1.6% and 2.5%, respectively. The prices of the assets at close of trading yesterday were \$20 and \$40 and the estimate of the coefficient of correlation between the returns on the two assets was 0.25. The parameter λ used in the EWMA model is 0.95.
- Calculate the current estimate of the covariance between the assets.
 - On the assumption that the prices of the assets at close of trading today are \$20.5 and \$40.5, update the correlation estimate.
- 22.10. The parameters of a GARCH(1,1) model are estimated as $\omega = 0.000004$, $\alpha = 0.05$, and $\beta = 0.92$. What is the long-run average volatility and what is the equation describing the way that the variance rate reverts to its long-run average? If the current volatility is 20% per year, what is the expected volatility in 20 days?
- 22.11. Suppose that the current daily volatilities of asset X and asset Y are 1.0% and 1.2%, respectively. The prices of the assets at close of trading yesterday were \$30 and \$50 and the estimate of the coefficient of correlation between the returns on the two assets made at this time was 0.50. Correlations and volatilities are updated using a GARCH(1,1) model. The estimates of the model's parameters are $\alpha = 0.04$ and $\beta = 0.94$. For the correlation $\omega = 0.000001$, and for the volatilities $\omega = 0.000003$. If the prices of the two assets at close of trading today are \$31 and \$51, how is the correlation estimate updated?

- 22.12. Suppose that the daily volatility of the FTSE 100 stock index (measured in pounds sterling) is 1.8% and the daily volatility of the dollar/sterling exchange rate is 0.9%. Suppose further that the correlation between the FTSE 100 and the dollar/sterling exchange rate is 0.4. What is the volatility of the FTSE 100 when it is translated to US dollars? Assume that the dollar/sterling exchange rate is expressed as the number of US dollars per pound sterling. (*Hint*: When $Z = XY$, the percentage daily change in Z is approximately equal to the percentage daily change in X plus the percentage daily change in Y .)
- 22.13. Suppose that in Problem 22.12 the correlation between the S&P 500 Index (measured in dollars) and the FTSE 100 Index (measured in sterling) is 0.7, the correlation between the S&P 500 Index (measured in dollars) and the dollar/sterling exchange rate is 0.3, and the daily volatility of the S&P 500 index is 1.6%. What is the correlation between the S&P 500 index (measured in dollars) and the FTSE 100 index when it is translated to dollars? (*Hint*: For three variables X , Y , and Z , the covariance between $X + Y$ and Z equals the covariance between X and Z plus the covariance between Y and Z .)
- 22.14. Show that the GARCH (1,1) model $\sigma_n^2 = \omega + \alpha u_{n-1}^2 + \beta \sigma_{n-1}^2$ in equation (22.9) is equivalent to the stochastic volatility model $dV = a(V_L - V)dt + \xi V dz$, where time is measured in days, V is the square of the volatility of the asset price, and

$$a = 1 - \alpha - \beta, \quad V_L = \frac{\omega}{1 - \alpha - \beta}, \quad \xi = \alpha\sqrt{2}$$

What is the stochastic volatility model when time is measured in years? (*Hint*: The variable u_{n-1} is the return on the asset price in time Δt . It can be assumed to be normally distributed with mean zero and standard deviation σ_{n-1} . It follows that the mean of u_{n-1}^2 and u_{n-1}^4 are σ_{n-1}^2 and $3\sigma_{n-1}^4$, respectively.)

- 22.15. At the end of Section 22.8, the VaR for the four-index example was calculated using the model-building approach. How does the VaR calculated change if the investment is \$2.5 million in each index? Carry out calculations when (a) volatilities and correlations are estimated using the equally weighted model and (b) when they are estimated using the EWMA model with $\lambda = 0.94$. Use the spreadsheets on the author's website.
- 22.16. What is the effect of changing λ from 0.94 to 0.97 in the EWMA calculations in the four-index example at the end of Section 22.8. Use the spreadsheets on the author's website.

Further Questions

- 22.17. Suppose that the price of gold at close of trading yesterday was \$600 and its volatility was estimated as 1.3% per day. The price at the close of trading today is \$596. Update the volatility estimate using
- The EWMA model with $\lambda = 0.94$
 - The GARCH(1,1) model with $\omega = 0.000002$, $\alpha = 0.04$, and $\beta = 0.94$.
- 22.18. Suppose that in Problem 22.17 the price of silver at the close of trading yesterday was \$16, its volatility was estimated as 1.5% per day, and its correlation with gold was estimated as 0.8. The price of silver at the close of trading today is unchanged at \$16. Update the volatility of silver and the correlation between silver and gold using the two models in Problem 22.17. In practice, is the ω parameter likely to be the same for gold and silver?

- 22.19. An Excel spreadsheet containing over 900 days of daily data on a number of different exchange rates and stock indices can be downloaded from the author's website:

www.rotman.utoronto.ca/~hull/data.

Choose one exchange rate and one stock index. Estimate the value of λ in the EWMA model that minimizes the value of $\sum_i (v_i - \beta_i)^2$, where v_i is the variance forecast made at the end of day $i - 1$ and β_i is the variance calculated from data between day i and day $i + 25$. Use the Solver tool in Excel. Set the variance forecast at the end of the first day equal to the square of the return on that day to start the EWMA calculations.

- 22.20. Suppose that the parameters in a GARCH (1,1) model are $\alpha = 0.03$, $\beta = 0.95$, and $\omega = 0.000002$.
- What is the long-run average volatility?
 - If the current volatility is 1.5% per day, what is your estimate of the volatility in 20, 40, and 60 days?
 - What volatility should be used to price 20-, 40-, and 60-day options?
 - Suppose that there is an event that increases the current volatility by 0.5% to 2% per day. Estimate the effect on the volatility in 20, 40, and 60 days.
 - Estimate by how much the event increases the volatilities used to price 20-, 40-, and 60-day options?
- 22.21. The calculations for the four-index example at the end of Section 22.8 assume that the investments in the DJIA, FTSE 100, CAC 40, and Nikkei 225 are \$4 million, \$3 million, \$1 million, and \$2 million, respectively. How does the VaR calculated change if the investments are \$3 million, \$3 million, \$1 million, and \$3 million, respectively? Carry out calculations when (a) volatilities and correlations are estimated using the equally weighted model and (b) when they are estimated using the EWMA model. What is the effect of changing λ from 0.94 to 0.90 in the EWMA calculations? Use the spreadsheets on the author's website.
- 22.22. Apply EWMA and GARCH(1,1) to data on the euro–USD exchange rate between July 27, 2005, and July 27, 2010. This data can be found on the author's website:

www.rotman.utoronto.ca/~hull/data.

23

CHAPTER



Credit Risk

The value-at-risk measure we covered in Chapter 21 and the Greek letters we studied in Chapter 18 are aimed at quantifying market risk. In this chapter we consider another important risk for financial institutions: credit risk. Most financial institutions devote considerable resources to the measurement and management of credit risk. Regulators have for many years required banks to keep capital to reflect the credit risks they are bearing. This capital is in addition to the capital, described in Business Snapshot 21.1, that is required for market risk.

Credit risk arises from the possibility that borrowers and counterparties in derivatives transactions may default. This chapter discusses a number of different approaches to estimating the probability that a company will default and explains the key difference between risk-neutral and real-world probabilities of default. It examines the nature of the credit risk in over-the-counter derivatives transactions and discusses the clauses derivatives dealers write into their contracts to reduce credit risk. It also covers default correlation, Gaussian copula models, and the estimation of credit value at risk.

Chapter 24 will discuss credit derivatives and show how ideas introduced in this chapter can be used to value these instruments.

23.1 CREDIT RATINGS

Rating agencies, such as Moody's, S&P, and Fitch, are in the business of providing ratings describing the creditworthiness of corporate bonds. The best rating assigned by Moody's is Aaa. Bonds with this rating are considered to have almost no chance of defaulting. The next best rating is Aa. Following that comes A, Baa, Ba, B, Caa, Ca, and C. Only bonds with ratings of Baa or above are considered to be *investment grade*. The S&P and Fitch ratings corresponding to Moody's Aaa, Aa, A, Baa, Ba, B, Caa, Ca, and C are AAA, AA, A, BBB, BB, B, CCC, CC, and C, respectively. To create finer rating measures, Moody's divides its Aa rating category into Aa1, Aa2, and Aa3, its A category into A1, A2, and A3, and so on. Similarly, S&P and Fitch divide their AA rating category into AA+, AA, and AA−, their A rating category into A+, A, and A−, and so on. Moody's Aaa category and the S&P/Fitch AAA category are not subdivided, nor usually are the two lowest rating categories.

23.2 HISTORICAL DEFAULT PROBABILITIES

Table 23.1 is typical of the data produced by rating agencies. It shows the default experience during a 20-year period of bonds that had a particular rating at the beginning of the period. For example, a bond with a credit rating of Baa has a 0.176% chance of defaulting by the end of the first year, a 0.494% chance of defaulting by the end of the second year, and so on. The probability of a bond defaulting during a particular year can be calculated from the table. For example, the probability that a bond initially rated Baa will default during the second year is $0.494 - 0.176 = 0.318\%$.

Table 23.1 shows that, for investment-grade bonds, the probability of default in a year tends to be an increasing function of time (e.g., the probabilities of an A-rated bond defaulting during years 0–5, 5–10, 10–15, and 15–20 are 0.717%, 1.329%, 1.526%, and 2.362%, respectively). This is because the bond issuer is initially considered to be creditworthy, and the more time that elapses, the greater the possibility that its financial health will decline. For bonds with a poor credit rating, the probability of default is often a decreasing function of time (e.g., the probabilities that a B-rated bond will default during years 0–5, 5–10, 10–15, and 15–20 are 25.895%, 18.482%, 11.721%, and 6.380%, respectively). The reason here is that, for a bond with a poor credit rating, the next year or two may be critical. The longer the issuer survives, the greater the chance that its financial health improves.

Hazard Rates

From Table 23.1 we can calculate the probability of a bond rated Caa or below defaulting during the third year as $38.682 - 29.384 = 9.298\%$. We will refer to this as the *unconditional default probability*. It is the probability of default during the third year as seen at time 0. The probability that the bond will survive until the end of year 2 is $100 - 29.384 = 70.616\%$. The probability that it will default during the third year conditional on no earlier default is therefore $0.09298/0.70616$, or 13.17%. Conditional default probabilities are referred to as *hazard rates* or *default intensities*.

The 13.17% we have just calculated is for a 1-year time period. Suppose instead that we consider a short time period of length Δt . The hazard rate $\lambda(t)$ at time t is then defined so that $\lambda(t) \Delta t$ is the probability of default between time t and $t + \Delta t$ conditional on no earlier default. If $V(t)$ is the cumulative probability of the company surviving to time t (i.e., no default by time t), the conditional probability of default between time t

Table 23.1 Average cumulative default rates (%), 1970–2009. *Source:* Moody's.

Term (years):	1	2	3	4	5	7	10	15	20
Aaa	0.000	0.012	0.012	0.037	0.105	0.245	0.497	0.927	1.102
Aa	0.022	0.059	0.091	0.159	0.234	0.384	0.542	1.150	2.465
A	0.051	0.165	0.341	0.520	0.717	1.179	2.046	3.572	5.934
Baa	0.176	0.494	0.912	1.404	1.926	2.996	4.851	8.751	12.327
Ba	1.166	3.186	5.583	8.123	10.397	14.318	19.964	29.703	37.173
B	4.546	10.426	16.188	21.256	25.895	34.473	44.377	56.098	62.478
Caa–C	17.723	29.384	38.682	46.094	52.286	59.771	71.376	77.545	80.211

and $t + \Delta t$ is $[V(t) - V(t + \Delta t)]/V(t)$. Since this equals $\lambda(t) \Delta t$, it follows that

$$V(t + \Delta t) - V(t) = -\lambda(t)V(t) \Delta t$$

Taking limits

$$\frac{dV(t)}{dt} = -\lambda(t)V(t)$$

from which

$$V(t) = e^{-\int_0^t \lambda(\tau) d\tau}$$

Defining $Q(t)$ as the probability of default by time t , so that $Q(t) = 1 - V(t)$, gives

$$Q(t) = 1 - e^{-\int_0^t \lambda(\tau) d\tau}$$

or

$$Q(t) = 1 - e^{-\bar{\lambda}(t)t} \quad (23.1)$$

where $\bar{\lambda}(t)$ is the average hazard rate (default intensity) between time 0 and time t .

23.3 RECOVERY RATES

When a company goes bankrupt, those that are owed money by the company file claims against the assets of the company.¹ Sometimes there is a reorganization in which these creditors agree to a partial payment of their claims. In other cases the assets are sold by the liquidator and the proceeds are used to meet the claims as far as possible. Some claims typically have priority over other claims and are met more fully.

The recovery rate for a bond is normally defined as the bond's market value a few days after a default, as a percent of its face value. Table 23.2 provides historical data on average recovery rates for different categories of bank loans and bonds in the United States. It shows that bank loans with a first lien on assets had the best average recovery rate, 65.6%. For bonds, the average recovery rate ranges from 49.8% for those that are

Table 23.2 Recovery rates on corporate bonds as a percentage of face value, 1982–2009. *Source:* Moody's.

<i>Class</i>	<i>Average recovery rate (%)</i>
First lien bank loan	65.6
Second lien bank loan	32.8
Senior unsecured bank loan	48.7
Senior secured bond	49.8
Senior unsecured bond	36.6
Senior subordinated bond	30.7
Subordinated bond	31.3
Junior subordinated bond	24.7

¹ In the United States, the claim made by a bond holder is the bond's face value plus accrued interest.

both senior to other lenders and secured to 24.7% for those that rank after other lenders with a security interest that is subordinate to other lenders.

Recovery rates are significantly negatively correlated with default rates.² This means that a bad year for the default rate is usually doubly bad because it is accompanied by a low recovery rate. For example, when the default rate for non-investment-grade bonds in a year is only 0.1%, the average recovery rate might be relatively high at 60%. When the default rate is relatively high at 3%, the average recovery rate might be only 35%.

23.4 ESTIMATING DEFAULT PROBABILITIES FROM BOND PRICES

The probability of default for a company can be estimated from the prices of bonds it has issued. The usual assumption is that the only reason a corporate bond sells for less than a similar risk-free bond is the possibility of default.³

Consider first an approximate calculation. Suppose that a bond yields 200 basis points more than a similar risk-free bond and that the expected recovery rate in the event of a default is 40%. The holder of a corporate bond must be expecting to lose 200 basis points (or 2% per year) from defaults. Given the recovery rate of 40%, this leads to an estimate of the probability of a default per year conditional on no earlier default of $0.02/(1 - 0.4)$, or 3.33%. In general,

$$\bar{\lambda} = \frac{s}{1 - R} \quad (23.2)$$

where $\bar{\lambda}$ is the average hazard rate (default intensity) per year, s is the spread of the corporate bond yield over the risk-free rate, and R is the expected recovery rate.

A More Exact Calculation

For a more exact calculation, suppose that the corporate bond we have been considering lasts for 5 years, provides a coupon 6% per annum (paid semiannually) and that the yield on the corporate bond is 7% per annum (with continuous compounding). The yield on a similar risk-free bond is 5% (with continuous compounding). The yields imply that the price of the corporate bond is 95.34 and the price of the risk-free bond is 104.09. The expected loss from default over the 5-year life of the bond is therefore $104.09 - 95.34$, or \$8.75. Suppose that the unconditional probability of default per year (assumed to be the same each year) is Q . Table 23.3 calculates the expected loss from default in terms of Q on the assumption that defaults can happen at times 0.5, 1.5, 2.5, 3.5, and 4.5 years (immediately before coupon payment dates). Risk-free rates for all maturities are assumed to be 5% (with continuous compounding).

To illustrate the calculations, consider the 3.5-year row in Table 23.3. The expected value of the corporate bond at time 3.5 years (calculated using forward interest rates

² See E.I. Altman, B. Brady, A. Resti, and A. Sironi, "The Link between Default and Recovery Rates: Theory, Empirical Evidence, and Implications," *Journal of Business*, 78, 6 (2005), 2203–28. This is also discussed in publications by Moody's Investors Service.

³ This assumption is not perfect. In practice the price of a corporate bond is affected by its liquidity. The lower the liquidity, the lower the price.

Table 23.3 Calculation of loss from default on a bond in terms of the default probabilities per year, Q . Notional principal = \$100.

<i>Time (years)</i>	<i>Default probability</i>	<i>Recovery amount (\$)</i>	<i>Risk-free value (\$)</i>	<i>Loss given default (\$)</i>	<i>Discount factor</i>	<i>PV of expected loss (\$)</i>
0.5	Q	40	106.73	66.73	0.9753	65.08 Q
1.5	Q	40	105.97	65.97	0.9277	61.20 Q
2.5	Q	40	105.17	65.17	0.8825	57.52 Q
3.5	Q	40	104.34	64.34	0.8395	54.01 Q
4.5	Q	40	103.46	63.46	0.7985	50.67 Q
<i>Total</i>						288.48 Q

and assuming no possibility of default) is

$$3 + 3e^{-0.05 \times 0.5} + 3e^{-0.05 \times 1.0} + 103e^{-0.05 \times 1.5} = 104.34$$

Given the definition of recovery rates in the previous section, the amount recovered if there is a default is 40, so that the loss given default is $104.34 - 40$, or \$64.34. The present value of this loss is 54.01. The expected loss is therefore 54.01 Q .

The total expected loss is 288.48 Q . Setting this equal to 8.75, we obtain a value for Q of $8.75/288.48$, or 3.03%. The calculations we have given assume that the default probability is the same in each year and that defaults take place at just one time during the year. We can extend the calculations to assume that defaults can take place more frequently. Also, instead of assuming a constant unconditional probability of default we can assume a constant hazard rate (default intensity) or assume a particular pattern for the variation of default probabilities with time. With several bonds we can estimate several parameters describing the term structure of default probabilities. Suppose, for example, we have bonds maturing in 3, 5, 7, and 10 years. We could use the first bond to estimate a default probability per year for the first 3 years, the second bond to estimate default probability per year for years 4 and 5, the third bond to estimate a default probability for years 6 and 7, and the fourth bond to estimate a default probability for years 8, 9, and 10 (see Problems 23.13 and 23.27). This approach is analogous to the bootstrap procedure in Section 4.5 for calculating a zero-coupon yield curve.

The Risk-Free Rate

A key issue when bond prices are used to estimate default probabilities is the meaning of the terms “risk-free rate” and “risk-free bond.” In equation (23.2), the spread s is the excess of the corporate bond yield over the yield on a similar risk-free bond. In Table 23.3, the risk-free value of the bond must be calculated using the risk-free discount rate. The benchmark risk-free rate that is usually used in quoting corporate bond yields is the yield on similar Treasury bonds. (For example, a bond trader might quote the yield on a particular corporate bond as being a spread of 250 basis points over Treasuries.)

As discussed in Section 4.1, traders usually use LIBOR/swap rates as proxies for risk-free rates when valuing derivatives. Traders also often use LIBOR/swap rates as risk-free rates when calculating default probabilities. For example, when they determine default probabilities from bond prices, the spread s in equation (23.2) is the spread of

the bond yield over the LIBOR/swap rate. Also, the risk-free discount rates used in the calculations in Table 23.3 are LIBOR/swap zero rates.

Credit default swaps (which will be discussed in the next chapter) can be used to imply the risk-free rate assumed by traders. The implied rate appears to be approximately equal to the LIBOR/swap rate minus 10 basis points on average.⁴ This estimate is plausible. As explained in Section 7.5, the credit risk in a swap rate is the credit risk from making a series of short-term loans to AA-rated counterparties and 10 basis points is a reasonable default risk premium for a AA-rated short-term instrument.

Asset Swaps

In practice, traders often use asset swap spreads as a way of extracting default probabilities from bond prices. This is because asset swap spreads provide a direct estimate of the spread of bond yields over the LIBOR/swap curve.

To explain how asset swaps work, consider the situation where an asset swap spread for a particular bond is quoted as 150 basis points. There are three possible situations:

1. The bond sells for its par value of 100. The swap then involves one side (company A) paying the coupon on the bond and the other side (company B) paying LIBOR plus 150 basis points. Note that it is the promised coupons that are exchanged. The exchanges take place regardless of whether the bond defaults.
2. The bond sells below its par value, say, for 95. The swap is then structured so that, in addition to the coupons, company A pays \$5 per \$100 of notional principal at the outset. Company B pays LIBOR plus 150 basis points.
3. The underlying bond sells above par, say, for 108. The swap is then structured so that, in addition to LIBOR plus 150 basis points, company B makes a payment of \$8 per \$100 of principal at the outset. Company A pays the coupons.

The effect of all this is that the present value of the asset swap spread is the amount by which the price of the corporate bond is exceeded by the price of a similar risk-free bond where the risk-free rate is assumed to be given by the LIBOR/swap curve (see Problem 23.22).

Consider again the example in Table 23.3 where the LIBOR/swap zero curve is flat at 5%. Suppose that instead of knowing the bond's price we know that the asset swap spread is 150 basis points. This means that the amount by which the value of the risk-free bond exceeds the value of the corporate bond is the present value of 150 basis points per year for 5 years. Assuming semiannual payments, this is \$6.55 per \$100 of principal. The total loss in Table 23.3 would in this case be set equal to \$6.55. This means that the default probability per year, Q , would be $6.55/288.48$, or 2.27%.

23.5 COMPARISON OF DEFAULT PROBABILITY ESTIMATES

The default probabilities estimated from historical data are usually much less than those derived from bond prices. The difference between the two was particularly large during

⁴ See J. Hull, M. Predescu, and A. White, "The Relationship between Credit Default Swap Spreads, Bond Yields, and Credit Rating Announcements," *Journal of Banking and Finance*, 28 (November 2004): 2789-2811.

the credit crisis which started in mid-2007. This is because there was what is termed a “flight to quality” during the crisis, where all investors wanted to hold safe securities such as Treasury bonds. The prices of corporate bonds declined, thereby increasing their yields. The credit spread s on these bonds increased and calculations such as the one in equation (23.2) gave very high default probability estimates.

We now show that it was also true that default probabilities calculated from bonds were higher than those calculated before the credit crisis. We first calculate the historical default probabilities using the data in the 7-year column of Table 23.1. (We use the 7-year column because the bonds we will look at later have a life of about 7 years.) From equation (23.1), we have

$$\bar{\lambda}(7) = -\frac{1}{7}\ln[1 - Q(7)]$$

where $\bar{\lambda}(t)$ is the average hazard rate (or default intensity) by time t and $Q(t)$ is the cumulative probability of default by time t . The values of $Q(7)$ for different rating categories are in Table 23.1. For example, for an A-rated company, $Q(7)$ is 0.01179. The average 7-year hazard rate is therefore

$$\bar{\lambda}(7) = -\frac{1}{7}\ln(1 - 0.01179) = 0.0017$$

or 0.17%.

To calculate average hazard rates from bond prices, we use equation (23.2) and bond yields published by Merrill Lynch. The results shown are averages between December 1996 and June 2007. The recovery rate is assumed to be 40%. The Merrill Lynch bonds have a life of about seven years. (This explains why we focused on the 7-year column in Table 23.1 when calculating historical default probabilities.) To calculate the bond yield spread, we assume, to be consistent with the arguments in the previous section, that the risk-free interest rate is the 7-year swap rate minus 10 basis points. For example, for A-rated bonds, the average Merrill Lynch yield was 5.995%. The average 7-year swap rate was 5.408%, so that the average risk-free rate was 5.308%. This gives the average 7-year hazard rate as

$$\frac{0.05995 - 0.05308}{1 - 0.4} = 0.0115$$

or 1.15%.

Table 23.4 shows that the ratio of the hazard rate backed out from bond prices to the hazard rate calculated from historical data is very high for investment-grade companies

Table 23.4 Seven-year average hazard rates (% per annum).

<i>Rating</i>	<i>Historical hazard rate</i>	<i>Hazard rate from bonds</i>	<i>Ratio</i>	<i>Difference</i>
Aaa	0.04	0.60	17.0	0.56
Aa	0.05	0.73	13.2	0.67
A	0.17	1.15	6.8	0.98
Baa	0.43	2.13	4.9	1.69
Ba	2.21	4.67	2.1	2.46
B	6.04	8.02	1.3	1.98
Caa and lower	13.01	18.39	1.4	5.39

Table 23.5 Expected excess return on bonds (basis points).

<i>Rating</i>	<i>Bond yield spread over Treasuries</i>	<i>Spread of risk-free rate over Treasuries</i>	<i>Spread for historical defaults</i>	<i>Excess return</i>
Aaa	78	42	2	34
Aa	86	42	3	41
A	111	42	10	59
Baa	169	42	26	101
Ba	322	42	132	148
B	523	42	362	119
Caa	1146	42	781	323

and tends to decline as a company's credit rating declines.⁵ The difference between the two hazard rates tends to increase as the credit rating declines.

Table 23.5 provides another way of looking at these results. It shows the excess return over the risk-free rate (still assumed to be the 7-year swap rate minus 10 basis points) earned by investors in bonds with different credit rating. Consider again an A-rated bond. The average spread over 7-year Treasuries is 111 basis points. Of this, 42 basis points are accounted for by the average spread between 7-year Treasuries and our proxy for the risk-free rate. A spread of 10 basis points is necessary to cover expected defaults. (This equals the historical hazard rate from Table 23.4 multiplied by 0.6 to allow for recoveries.) This leaves an excess return (after expected defaults have been taken into account) of 59 basis points.

Tables 23.4 and 23.5 show that a large percentage difference between default probability estimates translates into a small (but significant) excess return on the bond. For Aaa-rated bonds, the ratio of the two hazard rates is 17.0, but the expected excess return is only 34 basis points. The excess return tends to increase as credit quality declines.⁶

The excess return in Table 23.5 does not remain constant through time. Credit spreads, and therefore excess returns, were high in 2001, 2002, and the first half of 2003. After that they were fairly low until the credit crisis.

Real-World vs. Risk-Neutral Probabilities

The default probabilities implied from bond yields are risk-neutral probabilities of default. To explain why this is so, consider the calculations of default probabilities in Table 23.3. The calculations assume that expected default losses can be discounted at the risk-free rate. The risk-neutral valuation principle shows that this is a valid procedure providing the expected losses are calculated in a risk-neutral world. This means that the default probability Q in Table 23.3 must be a risk-neutral probability.

By contrast, the default probabilities implied from historical data are real-world default probabilities (sometimes also called *physical probabilities*). The expected excess return in Table 23.5 arises directly from the difference between real-world and risk-neutral default

⁵ The results in Tables 23.4 and 23.5 are updates of the results in J. Hull, M. Predescu, and A. White, "Bond Prices, Default Probabilities, and Risk Premiums," *Journal of Credit Risk*, 1, 2 (Spring 2005): 53–60.

⁶ The results for B-rated bonds in Tables 23.4 and 23.5 run counter to the overall pattern.

probabilities. If there were no expected excess return, then the real-world and risk-neutral default probabilities would be the same, and vice versa.

Why do we see such big differences between real-world and risk-neutral default probabilities? As we have just argued, this is the same as asking why corporate bond traders earn more than the risk-free rate on average.

One reason often advanced for the results is that corporate bonds are relatively illiquid and the returns on bonds are higher than they would otherwise be to compensate for this. This is true, but research shows that it does not fully explain the results in Table 23.5.⁷ Another possible reason for the results is that the subjective default probabilities of bond traders may be much higher than the those given in Table 23.1. Bond traders may be allowing for depression scenarios much worse than anything seen during the period covered by historical data. However, it is difficult to see how this can explain a large part of the excess return that is observed.

By far the most important reason for the results in Tables 23.4 and 23.5 is that bonds do not default independently of each other. There are periods of time when default rates are very low and periods of time when they are very high. Evidence for this can be obtained by looking at the default rates in different years. Moody's statistics show that between 1970 and 2009 the default rate per year ranged from a low of 0.09% in 1979 to highs of 3.97% and 5.35% in 2001 and 2009, respectively. The year-to-year variation in default rates gives rise to systematic risk (i.e., risk that cannot be diversified away) and bond traders earn an excess expected return for bearing the risk. (This is similar to the excess expected return earned by equity holders that is calculated by the capital asset pricing model—see the appendix to Chapter 3.) The variation in default rates from year to year may be because of overall economic conditions and it may be because a default by one company has a ripple effect resulting in defaults by other companies. (The latter is referred to by researchers as *credit contagion*.)

In addition to the systematic risk we have just talked about there is nonsystematic (or idiosyncratic) risk associated with each bond. If we were talking about stocks, we would argue that investors can diversify the nonsystematic risk by choosing a portfolio of, say, 30 stocks. They should not therefore demand a risk premium for bearing nonsystematic risk. For bonds, the arguments are not so clear-cut. Bond returns are highly skewed with limited upside. (For example, on an individual bond, there might be a 99.75% chance of a 7% return in a year, and a 0.25% chance of a –60% return in the year, the first outcome corresponding to no default and the second to default.) This type of risk is difficult to “diversify away”.⁸ It would require tens of thousands of different bonds. In practice, many bond portfolios are far from fully diversified. As a result, bond traders may earn an extra return for bearing nonsystematic risk as well as for bearing the systematic risk mentioned in the previous paragraph.

Which Default Probability Estimate Should Be Used?

At this stage it is natural to ask whether we should use real-world or risk-neutral default probabilities in the analysis of credit risk. The answer depends on the purpose of the

⁷ For example, J. Dick-Nielsen, P. Feldhütter, and D. Lando, “Corporate Bond Liquidity before and after the Onset of the Subprime Crisis,” Working Paper, Copenhagen Business School, 2010, uses a number of different liquidity measures and a large database of bond trades. It shows that the liquidity component of credit spreads is relatively small.

⁸ See J. D. Amato and E. M. Remolona, “The Credit Spread Puzzle,” *BIS Quarterly Review*, 5 (Dec. 2003): 51–63.

analysis. When valuing credit derivatives or estimating the impact of default risk on the pricing of instruments, risk-neutral default probabilities should be used. This is because the analysis calculates the present value of expected future cash flows and almost invariably (implicitly or explicitly) involves using risk-neutral valuation. When carrying out scenario analyses to calculate potential future losses from defaults, real-world default probabilities should be used.

23.6 USING EQUITY PRICES TO ESTIMATE DEFAULT PROBABILITIES

When we use a table such as Table 23.1 to estimate a company's real-world probability of default, we are relying on the company's credit rating. Unfortunately, credit ratings are revised relatively infrequently. This has led some analysts to argue that equity prices can provide more up-to-date information for estimating default probabilities.

In 1974, Merton proposed a model where a company's equity is an option on the assets of the company.⁹ Suppose, for simplicity, that a firm has one zero-coupon bond outstanding and that the bond matures at time T . Define:

- V_0 : Value of company's assets today
- V_T : Value of company's assets at time T
- E_0 : Value of company's equity today
- E_T : Value of company's equity at time T
- D : Debt repayment due at time T
- σ_V : Volatility of assets (assumed constant)
- σ_E : Instantaneous volatility of equity.

If $V_T < D$, it is (at least in theory) rational for the company to default on the debt at time T . The value of the equity is then zero. If $V_T > D$, the company should make the debt repayment at time T and the value of the equity at this time is $V_T - D$. Merton's model, therefore, gives the value of the firm's equity at time T as

$$E_T = \max(V_T - D, 0)$$

This shows that the equity is a call option on the value of the assets with a strike price equal to the repayment required on the debt. The Black–Scholes–Merton formula gives the value of the equity today as

$$E_0 = V_0 N(d_1) - De^{-rT} N(d_2) \quad (23.3)$$

where

$$d_1 = \frac{\ln(V_0/D) + (r + \sigma_V^2/2)T}{\sigma_V \sqrt{T}} \quad \text{and} \quad d_2 = d_1 - \sigma_V \sqrt{T}$$

The value of the debt today is $V_0 - E_0$.

The risk-neutral probability that the company will default on the debt is $N(-d_2)$. To calculate this, we require V_0 and σ_V . Neither of these are directly observable. However, if the company is publicly traded, we can observe E_0 . This means that equation (23.3) provides one condition that must be satisfied by V_0 and σ_V . We can also estimate σ_E

⁹ See R. Merton "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance*, 29 (1974): 449–70.

from historical data or options. From Itô's lemma,

$$\sigma_E E_0 = \frac{\partial E}{\partial V} \sigma_V V_0$$

or

$$\sigma_E E_0 = N(d_1) \sigma_V V_0 \quad (23.4)$$

This provides another equation that must be satisfied by V_0 and σ_V . Equations (23.3) and (23.4) provide a pair of simultaneous equations that can be solved for V_0 and σ_V .¹⁰

Example 23.1

The value of a company's equity is \$3 million and the volatility of the equity is 80%. The debt that will have to be paid in 1 year is \$10 million. The risk-free rate is 5% per annum. In this case $E_0 = 3$, $\sigma_E = 0.80$, $r = 0.05$, $T = 1$, and $D = 10$. Solving equations (23.3) and (23.4) yields $V_0 = 12.40$ and $\sigma_V = 0.2123$. The parameter d_2 is 1.1408, so that the probability of default is $N(-d_2) = 0.127$, or 12.7%. The market value of the debt is $V_0 - E_0$, or 9.40. The present value of the promised payment on the debt is $10e^{-0.05 \times 1} = 9.51$. The expected loss on the debt is therefore $(9.51 - 9.40)/9.51$, or about 1.2% of its no-default value. The expected loss (EL) equals the probability of default (PD) times one minus the recovery rate. It follows that the recovery rate equals one minus EL/PD. In this case, the recovery rate is $1 - 1.2/12.7$, or about 91%, of the debt's no-default value.

The basic Merton model we have just presented has been extended in a number of ways. For example, one version of the model assumes that a default occurs whenever the value of the assets falls below a barrier level. Another allows payments on debt instruments to be required at more than one time.

How well do the default probabilities produced by Merton's model and its extensions correspond to actual default experience? The answer is that Merton's model and its extensions produce a good ranking of default probabilities (risk-neutral or real-world). This means that a monotonic transformation can be used to convert the probability of default output from Merton's model into a good estimate of either the real-world or risk-neutral default probability.¹¹ It may seem strange to take a default probability $N(-d_2)$ that is in theory a risk-neutral default probability (because it is calculated from an option-pricing model) and use it to estimate a real-world default probability. Given the nature of the calibration process we have just described, the underlying assumption is that the ranking of the risk-neutral default probabilities of different companies is the same as the ranking of their real-world default probabilities.

23.7 CREDIT RISK IN DERIVATIVES TRANSACTIONS

The credit exposure on a derivatives transaction is more complicated than that on a loan. This is because the claim that will be made in the event of a default is more uncertain. Consider a financial institution that has one derivatives contract outstanding

¹⁰ To solve two nonlinear equations of the form $F(x, y) = 0$ and $G(x, y) = 0$, the Solver routine in Excel can be asked to find the values of x and y that minimize $[F(x, y)]^2 + [G(x, y)]^2$.

¹¹ Moody's KMV provides a service that transforms a default probability produced by Merton's model into a real-world default probability (which it refers to as an expected default frequency, or EDF). CreditGrades use Merton's model to estimate credit spreads, which are closely linked to risk-neutral default probabilities.

with a counterparty. Three possible situations can be distinguished:

1. Contract is always a liability to the financial institution
2. Contract is always an asset to the financial institution
3. Contract can become either an asset or a liability to the financial institution.

An example of a derivatives contract in the first category is a short option position; an example in the second category is a long option position; an example in the third category is a forward contract.

Derivatives in the first category have no credit risk to the financial institution. If the counterparty goes bankrupt, there will be no loss. The derivative is one of the counterparty's assets. It is likely to be retained, closed out, or sold to a third party. The result is no loss (or gain) to the financial institution.

Derivatives in the second category always have credit risk to the financial institution. If the counterparty goes bankrupt, a loss is likely to be experienced. The derivative is one of the counterparty's liabilities. The financial institution has to make a claim against the assets of the counterparty and may receive some percentage of the value of the derivative. (Typically, a claim arising from a derivatives transaction is unsecured and junior.)

Derivatives in the third category may or may not have credit risk. If the counterparty defaults when the value of the derivative is positive to the financial institution, a claim will be made against the assets of the counterparty and a loss is likely to be experienced. If the counterparty defaults when the value is negative to the financial institution, no loss is made because the derivative is retained, closed out, or sold to a third party.¹²

Adjusting Derivatives' Valuations for Counterparty Default Risk

How should a financial institution (or end-user of derivatives) adjust the value of a derivative to allow for counterparty credit risk? Consider a derivative that lasts until time T and has a value of f_0 today assuming no defaults. Let us suppose that defaults can take place at times t_1, t_2, \dots, t_n , where $t_n = T$, and that the value of the derivative to the financial institution (assuming no defaults) at time t_i is f_i . Define the risk-neutral probability of default at time t_i as q_i and the expected recovery rate as R .¹³

The exposure at time t_i is the financial institution's potential loss. This is $\max(f_i, 0)$. Assume that the expected recovery in the event of a default is R times the exposure. Assume also that the recovery rate and the probability of default are independent of the value of the derivative. The risk-neutral expected loss from default at time t_i is

$$q_i(1 - R)\hat{E}[\max(f_i, 0)]$$

where \hat{E} denotes expected value in a risk-neutral world. Taking present values leads to what is termed the *credit value adjustment (CVA)*:

$$\sum_{i=1}^n u_i v_i \tag{23.5}$$

where u_i equals $q_i(1 - R)$ and v_i is the value today of an instrument that pays off the exposure on the derivative under consideration at time t_i .

¹² Note that a company usually defaults because of a deterioration in its overall financial health, not because of the value of any one transaction.

¹³ The probability of default could be calculated from bond prices in the way described in Section 23.4.

Consider again the three categories of derivatives mentioned earlier. The first category (where the derivative is always a liability to the financial institution) is easy to deal with. The value of f_i is always negative and so the total expected loss from defaults given by equation (23.5) is always zero. The financial institution needs to make no adjustments for the cost of defaults. (Of course, the counterparty may want to take account of the possibility of the financial institution defaulting in its own pricing.)

For the second category (where the derivative is always an asset to the financial institution), f_i is always positive. This means that the expression $\max(f_i, 0)$ always equals f_i . Suppose that the only payoff from the derivative is at time T , the end of its life. In this case, f_0 must be the present value of f_i , so that $v_i = f_0$ for all i . The expression in equation (23.5) for the present value of the cost of defaults becomes

$$f_0 \sum_{i=1}^n q_i(1 - R)$$

If f_0^* is the actual value of the derivative (after allowing for possible defaults), then

$$f_0^* = f_0 - f_0 \sum_{i=1}^n q_i(1 - R) = f_0 \left[1 - \sum_{i=1}^n q_i(1 - R) \right] \quad (23.6)$$

One particular instrument that falls into the second category we are considering is an unsecured zero-coupon bond that promises \$1 at time T and is issued by the counterparty in the derivatives transaction. Define B_0 as the value of the bond assuming no possibility of default and B_0^* as the actual value of the bond. If we make the simplifying assumption that the recovery on the bond as a percent of its no-default value is the same as that on the derivative, then

$$B_0^* = B_0 \left[1 - \sum_{i=1}^n q_i(1 - R) \right] \quad (23.7)$$

From equations (23.6) and (23.7),

$$\frac{f_0^*}{f_0} = \frac{B_0^*}{B_0} \quad (23.8)$$

If y is the yield on a risk-free zero-coupon bond maturing at time T and y^* is the yield on a zero-coupon bond issued by the counterparty that matures at time T , then $B_0 = e^{-yT}$ and $B_0^* = e^{-y^*T}$, so that equation (23.8) gives

$$f_0^* = f_0 e^{-(y^* - y)T} \quad (23.9)$$

This shows that any derivative promising a payoff at time T can be valued by increasing the discount rate that is applied to the expected payoff in a risk-neutral world from the risk-free rate y to the risky rate y^* .

Example 23.2

Consider a 2-year over-the-counter option sold by company X with a value, assuming no possibility of default, of \$3. Suppose that 2-year zero-coupon bonds issued by the company X have a yield that is 1.5% greater than a similar risk-free zero-coupon bond. The value of the option is $3e^{-0.015 \times 2} = 2.91$, or \$2.91.

For the third category of derivatives, the sign of f_i is uncertain. The variable v_i is a call option on f_i with a strike price of zero. One way of calculating v_i is to simulate the underlying market variables over the life of the derivative. Sometimes approximate analytic calculations are possible (see, e.g., Problems 23.15 and 23.16).

The analyses we have presented assume that the probability of default is independent of the value of the derivative. This is likely to be a reasonable approximation in circumstances when the derivative is a small part of the portfolio of the counterparty or when the counterparty is using the derivative for hedging purposes. When a counterparty wants to enter into a large derivatives transaction for speculative purposes a financial institution should be wary. When the transaction has a large negative value for the counterparty (and a large positive value for the financial institution), the chance of counterparty declaring bankruptcy may be much higher than when the situation is the other way round.

Traders working for a financial institution use the term *right-way risk* to describe the situation where a counterparty is most likely to default when the financial institution has zero, or very little, exposure. They use the term *wrong-way risk* to describe the situation where the counterparty is most likely to default when the financial institution has a big exposure.

23.8 CREDIT RISK MITIGATION

In many instances the analysis we just have presented overstates the credit risk in a derivatives transaction. This is because there are a number of clauses that derivatives dealers include in their contracts to mitigate credit risk.

Netting

A clause that has become standard in the Master Agreements that govern transactions in the over-the-counter market is known as *netting*. This states that, if a company defaults on one transaction it has with a counterparty, it must default on all outstanding transactions with the counterparty.

Netting has been successfully tested in the courts in most jurisdictions. It can substantially reduce credit risk for a financial institution. Consider, for example, a financial institution that has three transactions outstanding with a particular counterparty. The transactions are worth +\$10 million, +\$30 million, and −\$25 million to the financial institution. Suppose the counterparty runs into financial difficulties and defaults on its outstanding obligations. To the counterparty, the three transactions have values of −\$10 million, −\$30 million, and +\$25 million, respectively. Without netting, the counterparty would default on the first two transactions and retain the third for a loss to the financial institution of \$40 million. With netting, it is compelled to default on all three transactions for a loss to the financial institution of \$15 million.¹⁴

Suppose a financial institution has a portfolio of N derivatives transactions with a particular counterparty. Suppose that the no-default value of the i th transaction is V_i

¹⁴ Note that, if the third transaction were worth −\$45 million to the financial institution instead of −\$25 million, the counterparty would choose not to default and there would be no loss to the financial institution.

and the amount recovered in the event of default is the recovery rate times this no default value. Without netting, the financial institution loses

$$(1 - R) \sum_{i=1}^N \max(V_i, 0)$$

where R is the recovery rate. With netting, it loses

$$(1 - R) \max\left(\sum_{i=1}^N V_i, 0\right)$$

Without netting, its loss is the payoff from a portfolio of call options on the transactions where each option has a strike price of zero. With netting, it is the payoff from a single option on the portfolio of transactions with a strike price of zero. The value of an option on a portfolio is never greater than, and is often considerably less than, the value of the corresponding portfolio of options.

The CVA analysis presented in the previous section can be extended so that equation (23.5) gives the present value of the expected loss from all transactions with a counterparty when netting agreements are in place. This is achieved by redefining v_i in the equation as the present value of a derivative that pays off the exposure at time t_i on the portfolio of all transactions with a counterparty.

A challenging task for a financial institution when considering whether it should enter into a new derivatives transaction with a counterparty is to calculate the incremental effect on expected credit losses. This can be done by using equation (23.5) in the way just described to calculate expected default costs with and without the transaction. It is interesting to note that, because of netting, the incremental effect of a new transaction on expected default losses can be negative. This happens when the value of the new transaction is negatively correlated with the value of existing transactions.

Collateralization

Another clause frequently used to mitigate credit risks is known as *collateralization*. Suppose that a company and a financial institution have entered into a number of derivatives transactions. A typical collateralization agreement specifies that the transactions be valued periodically. If the total value of the transactions to the financial institution is above a specified threshold level, the agreement requires the cumulative collateral posted by the company to equal the difference between the value of the transactions to the financial institution and the threshold level. If, after the collateral has been posted, the value of the transactions moves in favor of the company so that the difference between value of the transactions to the financial institution and the threshold level is less than the total margin already posted, the company can reclaim margin. In the event of a default by the company, the financial institution can seize the collateral. If the company does not post collateral as required, the financial institution can close out the transactions.

Suppose, for example, that the threshold level for the company is \$10 million and the transactions are marked to market daily for the purposes of collateralization. If on a particular day the value of the transactions to the financial institution rises from \$9 million to \$10.5 million, it can ask for \$0.5 million of collateral. If the next day the

value of the transactions rises further to \$11.4 million, it can ask for a further \$0.9 million of collateral. If the value of the transactions falls to \$10.9 million on the following day, the company can ask for \$0.5 million of the collateral to be returned. Note that the threshold (\$10 million in this case) can be regarded as a line of credit that the financial institution is prepared to grant to the company.

The margin must be deposited by the company with the financial institution in cash or in the form of acceptable securities such as bonds. The securities are subject to a discount known as a *haircut* applied to their market value for the purposes of margin calculations. Interest is normally paid on cash.

If the collateralization agreement is a two-way agreement a threshold will also be specified for the financial institution. The company can then ask the financial institution to post collateral when the value of the outstanding contracts to the company exceeds the threshold.

Collateralization agreements provide a great deal of protection against the possibility of default (just as the margin accounts discussed in Chapter 2 provide protection for people who trade futures on an exchange). However, the threshold amount is not subject to protection. Furthermore, even when the threshold is zero, the protection is not total. This is because, when a company gets into financial difficulties, it is likely to stop responding to requests to post collateral. By the time the counterparty exercises its right to close out contracts, their value may have moved further in its favor.

As explained in Chapter 2, over-the-counter derivatives are increasingly moving to clearing houses where market participants post both an initial margin and maintenance margins.

Downgrade Triggers

Another credit risk mitigation technique sometimes used by a financial institution is known as a *downgrade trigger*. This is a clause stating that if the credit rating of the counterparty falls below a certain level, say Baa, the financial institution has the option to close out a derivatives transaction at its market value.

Downgrade triggers do not provide protection from a big jump in a company's credit rating (for example, from A to default). Also, downgrade triggers work well only if relatively little use is made of them. If a company has many downgrade triggers outstanding with its counterparties, they are liable to provide little protection to any of the counterparties (see Business Snapshot 23.1).

23.9 DEFAULT CORRELATION

The term *default correlation* is used to describe the tendency for two companies to default at about the same time. There are a number of reasons why default correlation exists. Companies in the same industry or the same geographic region tend to be affected similarly by external events and as a result may experience financial difficulties at the same time. Economic conditions generally cause average default rates to be higher in some years than in other years. A default by one company may cause a default by another—the credit contagion effect. Default correlation means that credit risk cannot be completely diversified away and is the major reason why risk-neutral default probabilities are greater than real-world default probabilities (see Section 23.5).

Business Snapshot 23.1 Downgrade Triggers and Enron's Bankruptcy

In December 2001, Enron, one of the largest companies in the United States, went bankrupt. Right up to the last few days, it had an investment-grade credit rating. The Moody's rating immediately prior to default was Baa3 and the S&P rating was BBB-. The default was, however, anticipated to some extent by the stock market because Enron's stock price fell sharply in the period leading up to the bankruptcy. The probability of default estimated by models such as the one described in Section 23.6 increased sharply during this period.

Enron had entered into a huge number of derivatives transactions with downgrade triggers. The downgrade triggers stated that, if its credit rating fell below investment grade (i.e., below Baa3/BBB-), its counterparties would have the option of closing out the transactions. Suppose that Enron had been downgraded to below investment grade in, say, October 2001. The transactions that counterparties would choose to close out would be those with negative values to Enron (and positive values to the counterparties). So, Enron would have been required to make huge cash payments to its counterparties. It would not have been able to do this and immediate bankruptcy would have resulted.

This example illustrates that downgrade triggers provide protection only when relatively little use is made of them. When a company enters into a huge number of contracts with downgrade triggers, they may actually cause a company to go bankrupt prematurely. In Enron's case, we could argue that it was going to go bankrupt anyway and accelerating the event by two months would not have done any harm. In fact, Enron did have a chance of survival in October 2001. Attempts were being made to work out a deal with another energy company, Dynegy, and so forcing bankruptcy in October 2001 was not in the interests of either creditors or shareholders.

The credit rating companies found themselves in a difficult position. If they downgraded Enron to recognize its deteriorating financial position, they were signing its death warrant. If they did not do so, there was a chance of Enron surviving.

Default correlation is important in the determination of probability distributions for default losses from a portfolio of exposures to different counterparties.¹⁵ Two types of default correlation models that have been suggested by researchers are referred to as *reduced form models* and *structural models*.

Reduced form models assume that the hazard rates for different companies follow stochastic processes and are correlated with macroeconomic variables. When the hazard rate for company A is high there is a tendency for the hazard rate for company B to be high. This induces a default correlation between the two companies.

Reduced form models are mathematically attractive and reflect the tendency for economic cycles to generate default correlations. Their main disadvantage is that the range of default correlations that can be achieved is limited. Even when there is a perfect correlation between the hazard rates of the two companies, the probability that they will both default during the same short period of time is usually very low. This is liable to be a problem in some circumstances. For example, when two companies operate in the same

¹⁵ A binomial correlation measure that has been used by rating agencies is described in Technical Note 26 at www.rotman.utoronto.ca/~hull/TechnicalNotes.

industry and the same country or when the financial health of one company is for some reason heavily dependent on the financial health of another company, a relatively high default correlation may be warranted. One approach to solving this problem is by extending the model so that the hazard rate exhibits large jumps.

Structural models are based on a model similar to Merton's model (see Section 23.6). A company defaults if the value of its assets is below a certain level. Default correlation between companies A and B is introduced into the model by assuming that the stochastic process followed by the assets of company A is correlated with the stochastic process followed by the assets of company B. Structural models have the advantage over reduced form models that the correlation can be made as high as desired. Their main disadvantage is that they are liable to be computationally quite slow.

The Gaussian Copula Model for Time to Default

A model that has become a popular practical tool is the Gaussian copula model for the time to default. It can be characterized as a simplified structural model. It assumes that all companies will default eventually and attempts to quantify the correlation between the probability distributions of the times to default for two or more different companies.

The model can be used in conjunction with either real-world or risk-neutral default probabilities. The left tail of the real-world probability distribution for the time to default of a company can be estimated from data produced by rating agencies such as that in Table 23.1. The left tail of the risk-neutral probability distribution of the time to default can be estimated from bond prices using the approach in Section 23.4.

Define t_1 as the time to default of company 1 and t_2 as the time to default of company 2. If the probability distributions of t_1 and t_2 were normal, we could assume that the joint probability distribution of t_1 and t_2 is bivariate normal. As it happens, the probability distribution of a company's time to default is not even approximately normal. This is where a Gaussian copula model comes in. We transform t_1 and t_2 into new variables x_1 and x_2 using

$$x_1 = N^{-1}[Q_1(t_1)], \quad x_2 = N^{-1}[Q_2(t_2)]$$

where Q_1 and Q_2 are the cumulative, probability distributions for t_1 and t_2 , respectively, and N^{-1} is the inverse of the cumulative normal distribution ($u = N^{-1}(v)$ when $v = N(u)$). These are "percentile-to-percentile" transformations. The 5-percentile point in the probability distribution for t_1 is transformed to $x_1 = -1.645$, which is the 5-percentile point in the standard normal distribution; the 10-percentile point in the probability distribution for t_1 is transformed to $x_1 = -1.282$, which is the 10-percentile point in the standard normal distribution, and so on. The t_2 -to- x_2 transformation is similar.

By construction, x_1 and x_2 have normal distributions with mean zero and unit standard deviation. The model assumes that the joint distribution of x_1 and x_2 is bivariate normal. This assumption is referred to as using a *Gaussian copula*. The assumption is convenient because it means that the joint probability distribution of t_1 and t_2 is fully defined by the cumulative default probability distributions Q_1 and Q_2 for t_1 and t_2 , together with a single correlation parameter.

The attraction of the Gaussian copula model is that it can be extended to many companies. Suppose that we are considering n companies and that t_i is the time to default

of the i th company. We transform each t_i into a new variable, x_i , that has a standard normal distribution. The transformation is the percentile-to-percentile transformation

$$x_i = N^{-1}[Q_i(t_i)]$$

where Q_i is the cumulative probability distribution for t_i . It is then assumed that the x_i are multivariate normal. The default correlation between t_i and t_j is measured as the correlation between x_i and x_j . This is referred to as the *copula correlation*.¹⁶

The Gaussian copula is a useful way of representing the correlation structure between variables that are not normally distributed. It allows the correlation structure of the variables to be estimated separately from their marginal (unconditional) distributions. Although the variables themselves are not multivariate normal, the approach assumes that after a transformation is applied to each variable they are multivariate normal.

Example 23.3

Suppose that we wish to simulate defaults during the next 5 years in 10 companies. The copula default correlations between each pair of companies is 0.2. For each company the cumulative probability of a default during the next 1, 2, 3, 4, 5 years is 1%, 3%, 6%, 10%, 15%, respectively. When a Gaussian copula is used we sample from a multivariate normal distribution to obtain the x_i ($1 \leq i \leq 10$) with the pairwise correlation between the x_i being 0.2. We then convert the x_i to t_i , a time to default. When the sample from the normal distribution is less than $N^{-1}(0.01) = -2.33$, a default takes place within the first year; when the sample is between -2.33 and $N^{-1}(0.03) = -1.88$, a default takes place during the second year; when the sample is between -1.88 and $N^{-1}(0.06) = -1.55$, a default takes place during the third year; when the sample is between -1.55 and $N^{-1}(0.10) = -1.28$, a default takes place during the fourth year; when the sample is between -1.28 and $N^{-1}(0.15) = -1.04$, a default takes place during the fifth year. When the sample is greater than -1.04 , there is no default during the 5 years.

A Factor-Based Correlation Structure

To avoid defining a different correlation between x_i and x_j for each pair of companies i and j in the Gaussian copula model, a one-factor model is often used. The assumption is that

$$x_i = a_i F + \sqrt{1 - a_i^2} Z_i \quad (23.10)$$

In this equation, F is a common factor affecting defaults for all companies and Z_i is a factor affecting only company i . The variable F and the variables Z_i have independent standard normal distributions. The a_i are constant parameters between -1 and $+1$. The correlation between x_i and x_j is $a_i a_j$.¹⁷

Suppose that the probability that company i will default by a particular time T is $Q_i(T)$. Under the Gaussian copula model, a default happens by time T when

¹⁶ As an approximation, the copula correlation between t_i and t_j is often assumed to be the correlation between the equity returns for companies i and j .

¹⁷ The parameter a_i is sometimes approximated as the correlation of company i 's equity returns with a well-diversified market index.

$N(x_i) < Q_i(T)$ or $x_i < N^{-1}[Q_i(T)]$. From equation (23.10), this condition is

$$a_i F + \sqrt{1 - a_i^2} Z_i < N^{-1}[Q_i(T)]$$

or

$$Z_i < \frac{N^{-1}[Q_i(T)] - a_i F}{\sqrt{1 - a_i^2}}$$

Conditional on the value of the factor F , the probability of default is therefore

$$Q_i(T | F) = N\left(\frac{N^{-1}[Q_i(T)] - a_i F}{\sqrt{1 - a_i^2}}\right) \quad (23.11)$$

A particular case of the one-factor Gaussian copula model is where the probability distributions of default are the same for all i and the correlation between x_i and x_j is the same for all i and j . Suppose that $Q_i(T) = Q(T)$ for all i and that the common correlation is ρ , so that $a_i = \sqrt{\rho}$ for all i . Equation (23.11) becomes

$$Q(T | F) = N\left(\frac{N^{-1}[Q(T)] - \sqrt{\rho} F}{\sqrt{1 - \rho}}\right) \quad (23.12)$$

23.10 CREDIT VaR

Credit value at risk can be defined analogously to the way value at risk is defined for market risks (see Chapter 21). For example, a credit VaR with a confidence level of 99.9% and a 1-year time horizon is the credit loss that we are 99.9% confident will not be exceeded over 1 year.

Consider a bank with a very large portfolio of similar loans. As an approximation, assume that the probability of default is the same for each loan and the correlation between each pair of loans is the same. When the Gaussian copula model for time to default is used, the right-hand side of equation (23.12) is to a good approximation equal to the percentage of defaults by time T as a function of F . The factor F has a standard normal distribution. We are $X\%$ certain that its value will be greater than $N^{-1}(1 - X) = -N^{-1}(X)$. We are therefore $X\%$ certain that the percentage of losses over T years on a large portfolio will be less than $V(X, T)$, where

$$V(X, T) = N\left(\frac{N^{-1}[Q(T)] + \sqrt{\rho} N^{-1}(X)}{\sqrt{1 - \rho}}\right) \quad (23.13)$$

This result was first produced by Vasicek.¹⁸ As in equation (23.12), $Q(T)$ is the probability of default by time T and ρ is the copula correlation between any pair of loans.

A rough estimate of the credit VaR when an $X\%$ confidence level is used and the time horizon is T is therefore $L(1 - R)V(X, T)$, where L is the size of the loan portfolio and R is the recovery rate. The contribution of a particular loan of size L_i to the credit VaR

¹⁸ See O. Vasicek, "Probability of Loss on a Loan Portfolio," Working Paper, KMV, 1987. Vasicek's results were published in *Risk* magazine in December 2002 under the title "Loan Portfolio Value".

is $L_i(1 - R)V(X, T)$. This model underlies some of the formulas that regulators use for credit risk capital.¹⁹

Example 23.4

Suppose that a bank has a total of \$100 million of retail exposures. The 1-year probability of default averages 2% and the recovery rate averages 60%. The copula correlation parameter is estimated as 0.1. In this case,

$$V(0.999, 1) = N\left(\frac{N^{-1}(0.02) + \sqrt{0.1} N^{-1}(0.999)}{\sqrt{1 - 0.1}}\right) = 0.128$$

showing that the 99.9% worst case default rate is 12.8%. The 1-year 99.9% credit VaR is therefore $100 \times 0.128 \times (1 - 0.6)$ or \$5.13 million.

CreditMetrics

Many banks have developed other procedures for calculating credit VaR for internal use. One popular approach is known as CreditMetrics. This involves estimating a probability distribution of credit losses by carrying out a Monte Carlo simulation of the credit rating changes of all counterparties. Suppose we are interested in determining the probability distribution of losses over a 1-year period. On each simulation trial, we sample to determine the credit rating changes and defaults of all counterparties during the year. We then revalue our outstanding contracts to determine the total of credit losses for the year. After a large number of simulation trials, a probability distribution for credit losses is obtained. This can be used to calculate credit VaR.

This approach is liable to be computationally quite time intensive. However, it has the advantage that credit losses are defined as those arising from credit downgrades as well as defaults. Also the impact of credit mitigation clauses such as those described in Section 23.8 can be approximately incorporated into the analysis.

Table 23.6 is typical of the historical data provided by rating agencies on credit rating changes and could be used as a basis for a CreditMetrics Monte Carlo simulation. It shows the percentage probability of a bond moving from one rating category to another during a 1-year period. For example, a bond that starts with an A credit rating has a 90.91% chance of still having an A rating at the end of 1 year. It has a 0.05% chance of defaulting during the year, a 0.09% chance of dropping to B, and so on.²⁰

In sampling to determine credit losses, the credit rating changes for different counterparties should not be assumed to be independent. A Gaussian copula model is typically used to construct a joint probability distribution of rating changes similarly to the way it is used in the model in the previous section to describe the joint probability distribution of times to default. The copula correlation between the rating transitions for two companies is usually set equal to the correlation between their equity returns using a factor model similar to that in Section 23.9.

As an illustration of the CreditMetrics approach suppose that we are simulating the rating change of a Aaa and a Baa company over a 1-year period using the transition

¹⁹ For more details, see J. Hull, *Risk Management and Financial Institutions*, 2nd edn. Upper Saddle River, Pearson, 2010.

²⁰ Technical Note 11 at www.rotman.utoronto.ca/~hull/TechnicalNotes explains how a table such as Table 23.6 can be used to calculate transition matrices for periods other than 1 year.

Table 23.6 One-year ratings transition matrix, 1970–2009, with probabilities expressed as percentages and adjustments for transitions to the WR (without rating) category. *Source:* Moody's.

Initial rating	Rating at year-end								
	Aaa	Aa	A	Baa	Ba	B	Caa	Ca–C	Default
Aaa	90.57	8.76	0.63	0.01	0.03	0.00	0.00	0.00	0.00
Aa	1.06	90.30	8.19	0.36	0.05	0.02	0.01	0.00	0.02
A	0.06	2.90	90.91	5.44	0.50	0.09	0.03	0.00	0.05
Baa	0.04	0.20	4.91	89.18	4.44	0.83	0.19	0.02	0.17
Ba	0.01	0.07	0.42	6.24	83.47	7.99	0.58	0.09	1.13
B	0.01	0.04	0.15	0.39	5.40	82.50	6.35	0.79	4.37
Caa	0.00	0.02	0.02	0.19	0.51	9.55	70.01	4.97	14.72
Ca–C	0.00	0.00	0.00	0.00	0.43	3.01	11.61	51.67	33.28
Default	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	100.00

matrix in Table 23.6. Suppose that the correlation between the equities of the two companies is 0.2. On each simulation trial, we would sample two variables x_A and x_B from normal distributions so that their correlation is 0.2. The variable x_A determines the new rating of the Aaa company and variable x_B determines the new rating of the Baa company. Since $N^{-1}(0.9057) = 1.3147$, the Aaa company stays Aaa if $x_A < 1.3147$; since $N^{-1}(0.9057 + 0.0876) = 2.4730$, it becomes Aa if $1.3147 \leq x_A < 2.4730$; since $N^{-1}(0.9057 + 0.0876 + 0.0063) = 3.3528$, it becomes A if $2.4730 \leq x_A < 3.3528$; and so on. Consider next the Baa company. Since $N^{-1}(0.0004) = -3.3528$, the Baa company becomes Aaa if $x_B < -3.3528$; since $N^{-1}(0.0004 + 0.0020) = -2.8202$, it becomes Aa if $-3.3528 \leq x_B < -2.8202$; since

$$N^{-1}(0.0004 + 0.0020 + 0.0491) = -1.6305$$

it becomes A if $-2.8202 \leq x_B < -1.6305$; and so on. The Aaa never defaults during the year. The Baa defaults when $x_B > N^{-1}(0.9983)$, that is when $x_B > 2.9290$.

SUMMARY

The probability that a company will default during a particular period of time in the future can be estimated from historical data, bond prices, or equity prices. The default probabilities calculated from bond prices are risk-neutral probabilities, whereas those calculated from historical data are real-world probabilities. Real-world probabilities should be used for scenario analysis and the calculation of credit VaR. Risk-neutral probabilities should be used for valuing credit-sensitive instruments. Risk-neutral default probabilities are often significantly higher than real-world default probabilities.

The expected loss experienced from a counterparty default is reduced by what is known as netting. This is a clause in most contracts written by a financial institution stating that, if a counterparty defaults on one contract it has with the financial institution, it must default on all contracts it has with the financial institution. Losses are also reduced by collateralization and downgrade triggers. Collateralization requires the counterparty to

post collateral and a downgrade trigger gives a financial institution the option to close out a contract if the credit rating of a counterparty falls below a specified level.

Credit VaR can be defined similarly to the way VaR is defined for market risk. One approach to calculating it is the Gaussian copula model of time to default. This is used by regulators in the calculation of capital for credit risk. Another popular approach for calculating credit VaR is CreditMetrics. This uses a Gaussian copula model for credit rating changes.

FURTHER READING

- Altman, E. I. "Measuring Corporate Bond Mortality and Performance," *Journal of Finance*, 44 (1989): 902–22.
- Altman, E. I., B. Brady, A. Resti, and A. Sironi. "The Link Between Default and Recovery Rates: Theory, Empirical Evidence, and Implications," *Journal of Business*, 78, 6 (2005), 2203–28.
- Duffie, D., and K. Singleton "Modeling Term Structures of Defaultable Bonds," *Review of Financial Studies*, 12 (1999): 687–720.
- Gregory, J. *Counterparty Credit Risk*. Chichester, UK: Wiley, 2010.
- Finger, C. C. "A Comparison of Stochastic Default Rate Models," *RiskMetrics Journal*, 1 (November 2000): 49–73.
- Hull, J., M. Predescu, and A. White. "Relationship between Credit Default Swap Spreads, Bond Yields, and Credit Rating Announcements," *Journal of Banking and Finance*, 28 (November 2004): 2789–2811.
- Kealhofer, S. "Quantifying Credit Risk I: Default Prediction," *Financial Analysts Journal*, 59, 1 (2003a): 30–44.
- Kealhofer, S. "Quantifying Credit Risk II: Debt Valuation," *Financial Analysts Journal*, 59, 3 (2003b): 78–92.
- Li, D. X. "On Default Correlation: A Copula Approach," *Journal of Fixed Income*, March 2000: 43–54.
- Merton, R. C. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance*, 29 (1974): 449–70.
- Vasicek, O. "Loan Portfolio Value," *Risk* (December 2002), 160–62.

Practice Questions (Answers in the Solutions Manual)

- 23.1. The spread between the yield on a 3-year corporate bond and the yield on a similar risk-free bond is 50 basis points. The recovery rate is 30%. Estimate the average hazard rate per year over the 3-year period.
- 23.2. Suppose that in Problem 23.1 the spread between the yield on a 5-year bond issued by the same company and the yield on a similar risk-free bond is 60 basis points. Assume the same recovery rate of 30%. Estimate the average hazard rate per year over the 5-year period. What do your results indicate about the average hazard rate in years 4 and 5?
- 23.3. Should researchers use real-world or risk-neutral default probabilities for (a) calculating credit value at risk and (b) adjusting the price of a derivative for defaults?
- 23.4. How are recovery rates usually defined?
- 23.5. Explain the difference between an unconditional default probability density and a hazard rate.

- 23.6. Verify (a) that the numbers in the second column of Table 23.4 are consistent with the numbers in Table 23.1 and (b) that the numbers in the fourth column of Table 23.5 are consistent with the numbers in Table 23.4 and a recovery rate of 40%.
- 23.7. Describe how netting works. A bank already has one transaction with a counterparty on its books. Explain why a new transaction by a bank with a counterparty can have the effect of increasing or reducing the bank's credit exposure to the counterparty.
- 23.8. What is meant by a "haircut" in a collateralization agreement. A company offers to post its own equity as collateral. How would you respond?
- 23.9. Explain the difference between the Gaussian copula model for the time to default and CreditMetrics as far as the following are concerned: (a) the definition of a credit loss and (b) the way in which default correlation is modeled.
- 23.10. Suppose that the LIBOR/swap curve is flat at 6% with continuous compounding and a 5-year bond with a coupon of 5% (paid semiannually) sells for 90.00. How would an asset swap on the bond be structured? What is the asset swap spread that would be calculated in this situation?
- 23.11. Show that the value of a coupon-bearing corporate bond is the sum of the values of its constituent zero-coupon bonds when the amount claimed in the event of default is the no-default value of the bond, but that this is not so when the claim amount is the face value of the bond plus accrued interest.
- 23.12. A 4-year corporate bond provides a coupon of 4% per year payable semiannually and has a yield of 5% expressed with continuous compounding. The risk-free yield curve is flat at 3% with continuous compounding. Assume that defaults can take place at the end of each year (immediately before a coupon or principal payment) and that the recovery rate is 30%. Estimate the risk-neutral default probability on the assumption that it is the same each year.
- 23.13. A company has issued 3- and 5-year bonds with a coupon of 4% per annum payable annually. The yields on the bonds (expressed with continuous compounding) are 4.5% and 4.75%, respectively. Risk-free rates are 3.5% with continuous compounding for all maturities. The recovery rate is 40%. Defaults can take place halfway through each year. The risk-neutral default rates per year are Q_1 for years 1 to 3 and Q_2 for years 4 and 5. Estimate Q_1 and Q_2 .
- 23.14. Suppose that a financial institution has entered into a swap dependent on the sterling interest rate with counterparty X and an exactly offsetting swap with counterparty Y. Which of the following statements are true and which are false?
- (a) The total present value of the cost of defaults is the sum of the present value of the cost of defaults on the contract with X plus the present value of the cost of defaults on the contract with Y.
 - (b) The expected exposure in 1 year on both contracts is the sum of the expected exposure on the contract with X and the expected exposure on the contract with Y.
 - (c) The 95% upper confidence limit for the exposure in 1 year on both contracts is the sum of the 95% upper confidence limit for the exposure in 1 year on the contract with X and the 95% upper confidence limit for the exposure in 1 year on the contract with Y.

Explain your answers.

- 23.15. A company enters into a 1-year forward contract to sell \$100 for AUD150. The contract is initially at the money. In other words, the forward exchange rate is 1.50. The 1-year dollar risk-free rate of interest is 5% per annum. The 1-year dollar rate of interest at which the counterparty can borrow is 6% per annum. The exchange rate volatility is 12% per annum. Estimate the present value of the cost of defaults on the contract. Assume that defaults are recognized only at the end of the life of the contract.
- 23.16. Suppose that in Problem 23.15, the 6-month forward rate is also 1.50 and the 6-month dollar risk-free interest rate is 5% per annum. Suppose further that the 6-month dollar rate of interest at which the counterparty can borrow is 5.5% per annum. Estimate the present value of the cost of defaults assuming that defaults can occur either at the 6-month point or at the 1-year point? (If a default occurs at the 6-month point, the company's potential loss is the market value of the contract.)
- 23.17. "A long forward contract subject to credit risk is a combination of a short position in a no-default put and a long position in a call subject to credit risk." Explain this statement.
- 23.18. Explain why the credit exposure on a matched pair of forward contracts resembles a straddle.
- 23.19. Explain why the impact of credit risk on a matched pair of interest rate swaps tends to be less than that on a matched pair of currency swaps.
- 23.20. "When a bank is negotiating currency swaps, it should try to ensure that it is receiving the lower interest rate currency from a company with a low credit risk." Explain why.
- 23.21. Does put–call parity hold when there is default risk? Explain your answer.
- 23.22. Suppose that in an asset swap B is the market price of the bond per dollar of principal, B^* is the default-free value of the bond per dollar of principal, and V is the present value of the asset swap spread per dollar of principal. Show that $V = B^* - B$.
- 23.23. Show that under Merton's model in Section 23.6 the credit spread on a T -year zero-coupon bond is $-\ln[N(d_2) + N(-d_1)/L]/T$, where $L = De^{-rT}/V_0$.
- 23.24. Suppose that the spread between the yield on a 3-year zero-coupon riskless bond and a 3-year zero-coupon bond issued by a corporation is 1%. By how much does Black–Scholes–Merton overstate the value of a 3-year European option sold by the corporation.
- 23.25. Give an example of (a) right-way risk and (b) wrong-way risk.

Further Questions

- 23.26. Suppose a 3-year corporate bond provides a coupon of 7% per year payable semi-annually and has a yield of 5% (expressed with semiannual compounding). The yields for all maturities on risk-free bonds is 4% per annum (expressed with semiannual compounding). Assume that defaults can take place every 6 months (immediately before a coupon payment) and the recovery rate is 45%. Estimate the default probabilities assuming (a) that the unconditional default probabilities are the same on each possible default date and (b) that the default probabilities conditional on no earlier default are the same on each possible default date.
- 23.27. A company has 1- and 2-year bonds outstanding, each providing a coupon of 8% per year payable annually. The yields on the bonds (expressed with continuous compounding) are

- 6.0% and 6.6%, respectively. Risk-free rates are 4.5% for all maturities. The recovery rate is 35%. Defaults can take place halfway through each year. Estimate the risk-neutral default rate each year.
- 23.28. Explain carefully the distinction between real-world and risk-neutral default probabilities. Which is higher? A bank enters into a credit derivative where it agrees to pay \$100 at the end of 1 year if a certain company's credit rating falls from A to Baa or lower during the year. The 1-year risk-free rate is 5%. Using Table 23.6, estimate a value for the derivative. What assumptions are you making? Do they tend to overstate or understate the value of the derivative.
- 23.29. The value of a company's equity is \$4 million and the volatility of its equity is 60%. The debt that will have to be repaid in 2 years is \$15 million. The risk-free interest rate is 6% per annum. Use Merton's model to estimate the expected loss from default, the probability of default, and the recovery rate in the event of default. (*Hint*: The Solver function in Excel can be used for this question, as indicated in footnote 10.)
- 23.30. Suppose that a bank has a total of \$10 million of exposures of a certain type. The 1-year probability of default averages 1% and the recovery rate averages 40%. The copula correlation parameter is 0.2. Estimate the 99.5% 1-year credit VaR.



24

CHAPTER

Credit Derivatives

An important development in derivatives markets since the late 1990s has been the growth of credit derivatives. In 2000, the total notional principal for outstanding credit derivatives contracts was about \$800 billion. By December 2009, this had become \$32 trillion. Credit derivatives are contracts where the payoff depends on the creditworthiness of one or more companies or countries. This chapter explains how credit derivatives work and how they are valued.

Credit derivatives allow companies to trade credit risks in much the same way that they trade market risks. Banks and other financial institutions used to be in the position where they could do little once they had assumed a credit risk except wait (and hope for the best). Now they can actively manage their portfolios of credit risks, keeping some and entering into credit derivatives contracts to protect themselves from others. As indicated in Business Snapshot 24.1, banks have been the biggest buyers of credit protection and insurance companies have been the biggest sellers.

Credit derivatives can be categorized as “single-name” or “multi-name.” The most popular single-name credit derivative is a credit default swap. The payoff from this instrument depends on the creditworthiness of one company or country. There are two sides to the contract: the buyer and seller of protection. There is a payoff from the seller of protection to the buyer of protection if the specified entity (company or country) defaults on its obligations. The most popular multi-name credit derivative is a collateralized debt obligation. In this, a portfolio of debt instruments is specified and a complex structure is created where the cash flows from the portfolio are channelled to different categories of investors. Chapter 8 describes how multi-name credit derivatives were created from residential mortgages during the period leading up to the credit crisis. This chapter focuses on the situation where the underlying credit risks are those of corporations or countries. Multi-name credit derivatives increased in popularity relative to single-name credit derivatives up to June 2007 but became less popular during the 2007–2009 credit crisis.

This chapter starts by explaining how credit default swaps work and how they are valued. It then covers the trading of forwards and options on credit default swaps and total return swaps. It explains credit indices, basket credit default swaps, asset-backed securities, and collateralized debt obligations. It expands on the material in Chapter 23 to show how the Gaussian copula model of default correlation can be used to value tranches of collateralized debt obligations.

Business Snapshot 24.1 Who Bears the Credit Risk?

Traditionally banks have been in the business of making loans and then bearing the credit risk that the borrower will default. However, banks have for some time been reluctant to keep loans on their balance sheets. This is because, after the capital required by regulators has been accounted for, the average return earned on loans is often less attractive than that on other assets. As discussed in Section 8.1, banks created asset-backed securities to pass loans (and their credit risk) on to investors. In the late 1990s and early 2000s, banks also made extensive use of credit derivatives to shift the credit risk in their loans to other parts of the financial system.

If banks have been net buyers of credit protection, who have been net sellers? The answer is insurance companies. Insurance companies are not regulated in the same way as banks and as a result are sometimes more willing to bear credit risks than banks.

The result of all this is that the financial institution bearing the credit risk of a loan is often different from the financial institution that did the original credit checks. As the credit crisis of 2007 has shown, this is not always good for the overall health of the financial system.

24.1 CREDIT DEFAULT SWAPS

The most popular credit derivative is a *credit default swap* (CDS). This is a contract that provides insurance against the risk of a default by particular company. The company is known as the *reference entity* and a default by the company is known as a *credit event*. The buyer of the insurance obtains the right to sell bonds issued by the company for their face value when a credit event occurs and the seller of the insurance agrees to buy the bonds for their face value when a credit event occurs.¹ The total face value of the bonds that can be sold is known as the credit default swap's *notional principal*.

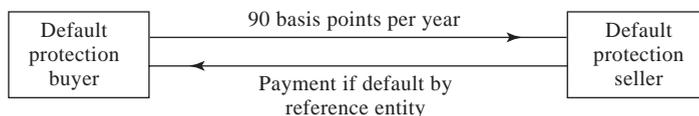
The buyer of the CDS makes periodic payments to the seller until the end of the life of the CDS or until a credit event occurs. These payments are typically made in arrears every quarter, but deals where payments are made every month, 6 months, or 12 months also occur and sometimes payments are made in advance. The settlement in the event of a default involves either physical delivery of the bonds or a cash payment.

An example will help to illustrate how a typical deal is structured. Suppose that two parties enter into a 5-year credit default swap on March 20, 2012. Assume that the notional principal is \$100 million and the buyer agrees to pay 90 basis points per annum for protection against default by the reference entity, with payments being made quarterly in arrears.

The CDS is shown in Figure 24.1. If the reference entity does not default (i.e., there is no credit event), the buyer receives no payoff and pays 22.5 basis points (a quarter of 90 basis points) on \$100 million on June 20, 2012, and every quarter thereafter until March 20, 2017. The amount paid each quarter is $0.00225 \times 100,000,000$, or \$225,000.²

¹ The face value (or par value) of a coupon-bearing bond is the principal amount that the issuer repays at maturity if it does not default.

² The quarterly payments are liable to be slightly different from \$225,000 because of the application of the day count conventions described in Chapter 6.

Figure 24.1 Credit default swap.

If there is a credit event, a substantial payoff is likely. Suppose that the buyer notifies the seller of a credit event on May 20, 2015 (2 months into the fourth year). If the contract specifies physical settlement, the buyer has the right to sell bonds issued by the reference entity with a face value of \$100 million for \$100 million. If, as is now usual, there is cash settlement, an ISDA-organized auction process is used to determine the mid-market value of the cheapest deliverable bond several days after the credit event. Suppose the auction indicates that the bond is worth \$35 per \$100 of face value. The cash payoff would be \$65 million.

The regular quarterly, semiannual, or annual payments from the buyer of protection to the seller of protection cease when there is a credit event. However, because these payments are made in arrears, a final accrual payment by the buyer is usually required. In our example, where there is a default on May 20, 2015, the buyer would be required to pay to the seller the amount of the annual payment accrued between March 20, 2015, and May 20, 2015 (approximately \$150,000), but no further payments would be required.

The total amount paid per year, as a percent of the notional principal, to buy protection (90 basis points in our example) is known as the *CDS spread*. Several large banks are market makers in the credit default swap market. When quoting on a new 5-year credit default swap on a company, a market maker might bid 250 basis points and offer 260 basis points. This means that the market maker is prepared to buy protection by paying 250 basis points per year (i.e., 2.5% of the principal per year) and to sell protection for 260 basis points per year (i.e., 2.6% of the principal per year).

Many different companies and countries are reference entities for the CDS contracts that trade. As mentioned, payments are usually made quarterly in arrears. Contracts with maturities of 5 years are most popular, but other maturities such as 1, 2, 3, 7, and 10 years are not uncommon. Usually contracts mature on one of the following standard dates: March 20, June 20, September 20, and December 20. The effect of this is that the actual time to maturity of a contract when it is initiated is close to, but not necessarily the same as, the number of years to maturity that is specified. Suppose you call a dealer on November 15, 2012, to buy 5-year protection on a company. The contract would probably last until December 20, 2017. Your first payment would be due on December 20, 2012, and would equal an amount covering the November 15, 2012, to December 20, 2012, period.³ A key aspect of a CDS contract is the definition of a credit event (i.e., a default). Usually a credit event is defined as a failure to make a payment as it becomes due, a restructuring of debt, or a bankruptcy. Restructuring is sometimes excluded in North American contracts, particularly in situations where the yield on the company's debt is high. More information on the CDS market is given in Business Snapshot 24.2.

³ If the time to the first standard date is less than 1 month, then the first payment is typically made on the second standard payment date; otherwise it is made on the first standard payment date.

Business Snapshot 24.2 The CDS Market

In 1998 and 1999, the International Swaps and Derivatives Association (ISDA) developed a standard contract for trading credit default swaps in the over-the-counter market. Since then the market has grown very fast. A CDS contract is like an insurance contract in many ways, but there is one key difference. An insurance contract provides protection against losses on an asset that is owned. In the case of a CDS, the underlying asset does not have to be owned.

During the credit turmoil that started in August 2007, regulators became very concerned about systemic risk (see Business Snapshot 2.3). They felt that credit default swaps were a source of vulnerability for financial markets. The danger is that a default by one financial institution might lead to big losses by its counterparties in CDS transactions and further defaults by other financial institutions. Regulatory concerns were fueled by troubles at insurance giant AIG. This was a big seller of protection on the AAA-rated tranches created from mortgages (see Chapter 8). The protection proved very costly to AIG and the company was bailed out by the U.S. government.

Regulatory concerns have led to the development of clearing houses for CDS trades (and other over-the-counter derivatives) between financial institutions. The clearing house requires market participants to post margin on their CDS trades in much the same way that traders post margin on futures contracts.

During 2007 and 2008, trading ceased in many types of credit derivatives, but CDSs continued to trade actively (although the cost of protection increased dramatically). The advantage of CDSs over some other credit derivatives is that the way they work is straightforward. Other credit derivatives, such as those created from the securitization of household mortgages (see Chapter 8), lack this transparency.

It is not uncommon for the volume of CDSs on a company to be greater than its debt. Cash settlement of contracts is then clearly necessary. When Lehman defaulted in September 2008, there was about \$400 billion of CDS contracts and \$155 billion of Lehman debt outstanding. The payout to the buyers of protection (determined by an auction process) was 91.375% of principal.

Credit Default Swaps and Bond Yields

A CDS can be used to hedge a position in a corporate bond. Suppose that an investor buys a 5-year corporate bond yielding 7% per year for its face value and at the same time enters into a 5-year CDS to buy protection against the issuer of the bond defaulting. Suppose that the CDS spread is 200 basis points, or 2%, per annum. The effect of the CDS is to convert the corporate bond to a risk-free bond (at least approximately). If the bond issuer does not default the investor earns 5% per year when the CDS spread is netted against the corporate bond yield. If the bond does default the investor earns 5% up to the time of the default. Under the terms of the CDS, the investor is then able to exchange the bond for its face value. This face value can be invested at the risk-free rate for the remainder of the 5 years.

This shows that the excess of an n -year bond yield over the risk-free rate should approximately equal the n -year CDS spread. If it is markedly more than this, an investor can earn more than the risk-free rate by buying the corporate bond and buying protection. If it is markedly less than this, an investor can borrow at less than

the risk-free rate by shorting the corporate bond and selling CDS protection. The relevant risk-free rate is usually assumed to be the LIBOR/swap rate, so that the excess of the bond yield over the risk-free rate is the asset swap spread (see Section 23.4).

The CDS–bond basis is defined as

$$\text{CDS–bond basis} = \text{CDS spread} - \text{Excess of bond yield over risk-free rate}$$

or equivalently

$$\text{CDS–bond basis} = \text{CDS spread} - \text{Asset swap spread}$$

The arbitrage argument given above suggests that this should be close to zero. Prior to the 2007 credit crisis, it was on average slightly positive. During the crisis, it tended to be negative and became highly negative for a short period of time in January 2009.

The Cheapest-to-Deliver Bond

As explained in Section 23.3, the recovery rate on a bond is defined as the value of the bond immediately after default as a percent of face value. This means that the payoff from a CDS is $L(1 - R)$, where L is the notional principal and R is the recovery rate.

Usually a CDS specifies that a number of different bonds can be delivered in the event of a default. The bonds typically have the same seniority, but they may not sell for the same percentage of face value immediately after a default.⁴ This gives the holder of a CDS a cheapest-to-deliver bond option. As already mentioned, an auction process, organized by ISDA, is usually used to determine the value of the cheapest-to-deliver bond and, therefore, the payoff to the buyer of protection.

24.2 VALUATION OF CREDIT DEFAULT SWAPS

The CDS spread for a particular reference entity can be calculated from default probability estimates. We will illustrate how this is done with a simple example.

Suppose that the probability of a reference entity defaulting during a year conditional on no earlier default is 2%.⁵ Table 24.1 shows survival probabilities and unconditional default probabilities (i.e., default probabilities as seen at time zero) for each of the 5 years. The probability of a default during the first year is 0.02 and the probability the reference entity will survive until the end of the first year is 0.98. The probability of a default during the second year is $0.02 \times 0.98 = 0.0196$ and the probability of survival until the end of the second year is $0.98 \times 0.98 = 0.9604$. The probability of default during the third year is $0.02 \times 0.9604 = 0.0192$, and so on.

We will assume that defaults always happen halfway through a year and that payments on the credit default swap are made once a year, at the end of each year. We also assume that the risk-free (LIBOR) interest rate is 5% per annum with continuous compounding

⁴ There are a number of reasons for this. The claim that is made in the event of a default is typically equal to the bond's face value plus accrued interest. Bonds with high accrued interest at the time of default therefore tend to have higher prices immediately after default. Also the market may judge that in the event of a reorganization of the company some bond holders will fare better than others.

⁵ This is a hazard rate expressed with annual compounding. The equivalent continuously compounded hazard rate is 2.02%.

Table 24.1 Unconditional default probabilities and survival probabilities.

<i>Time (years)</i>	<i>Default probability</i>	<i>Survival probability</i>
1	0.0200	0.9800
2	0.0196	0.9604
3	0.0192	0.9412
4	0.0188	0.9224
5	0.0184	0.9039

and the recovery rate is 40%. There are three parts to the calculation. These are shown in Tables 24.2, 24.3, and 24.4.

Table 24.2 shows the calculation of the present value of the expected payments made on the CDS assuming that payments are made at the rate of s per year and the notional principal is \$1. For example, there is a 0.9412 probability that the third payment of s is made. The expected payment is therefore $0.9412s$ and its present value is $0.9412se^{-0.05 \times 3} = 0.8101s$. The total present value of the expected payments is $4.0704s$.

Table 24.3 shows the calculation of the present value of the expected payoff assuming a notional principal of \$1. As mentioned earlier, we are assuming that defaults always happen halfway through a year. For example, there is a 0.0192 probability of a payoff halfway through the third year. Given that the recovery rate is 40%, the expected payoff at this time is $0.0192 \times 0.6 \times 1 = 0.0115$. The present value of the expected payoff is $0.0115e^{-0.05 \times 2.5} = 0.0102$. The total present value of the expected payoffs is \$0.0511.

As a final step, Table 24.4 considers the accrual payment made in the event of a default. For example, there is a 0.0192 probability that there will be a final accrual payment halfway through the third year. The accrual payment is $0.5s$. The expected accrual payment at this time is therefore $0.0192 \times 0.5s = 0.0096s$. Its present value is $0.0096se^{-0.05 \times 2.5} = 0.0085s$. The total present value of the expected accrual payments is $0.0426s$.

Table 24.2 Calculation of the present value of expected payments.
Payment = s per annum.

<i>Time (years)</i>	<i>Probability of survival</i>	<i>Expected payment</i>	<i>Discount factor</i>	<i>PV of expected payment</i>
1	0.9800	$0.9800s$	0.9512	$0.9322s$
2	0.9604	$0.9604s$	0.9048	$0.8690s$
3	0.9412	$0.9412s$	0.8607	$0.8101s$
4	0.9224	$0.9224s$	0.8187	$0.7552s$
5	0.9039	$0.9039s$	0.7788	$0.7040s$
<i>Total</i>				$4.0704s$

Table 24.3 Calculation of the present value of expected payoff.

Notional principal = \$1.

<i>Time (years)</i>	<i>Probability of default</i>	<i>Recovery rate</i>	<i>Expected payoff (\$)</i>	<i>Discount factor</i>	<i>PV of expected payoff (\$)</i>
0.5	0.0200	0.4	0.0120	0.9753	0.0117
1.5	0.0196	0.4	0.0118	0.9277	0.0109
2.5	0.0192	0.4	0.0115	0.8825	0.0102
3.5	0.0188	0.4	0.0113	0.8395	0.0095
4.5	0.0184	0.4	0.0111	0.7985	0.0088
<i>Total</i>					0.0511

From Tables 24.2 and 24.4, the present value of the expected payments is

$$4.0704s + 0.0426s = 4.1130s$$

From Table 24.3, the present value of the expected payoff is 0.0511. Equating the two gives

$$4.1130s = 0.0511$$

or $s = 0.0124$. The mid-market CDS spread for the 5-year deal we have considered should be 0.0124 times the principal or 124 basis points per year. This result can also be produced using the DerivaGem CDS worksheet. The hazard rate (continuously compounded in DerivaGem) should be input as 2.02% for all maturities, the term structure is flat at 5%, and the recovery rate is 40%.

The calculations assume that defaults happen only at points midway between payment dates. This simple assumption can be relaxed, but usually gives good results.

Marking to Market a CDS

A CDS, like most other swaps, is marked to market daily. It may have a positive or negative value. Suppose, for example the credit default swap in our example had been negotiated some time ago for a spread of 150 basis points, the present value of the payments by the buyer would be $4.1130 \times 0.0150 = 0.0617$ and the present value of the

Table 24.4 Calculation of the present value of accrual payment.

<i>Time (years)</i>	<i>Probability of default</i>	<i>Expected accrual payment</i>	<i>Discount factor</i>	<i>PV of expected accrual payment</i>
0.5	0.0200	0.0100s	0.9753	0.0097s
1.5	0.0196	0.0098s	0.9277	0.0091s
2.5	0.0192	0.0096s	0.8825	0.0085s
3.5	0.0188	0.0094s	0.8395	0.0079s
4.5	0.0184	0.0092s	0.7985	0.0074s
<i>Total</i>				0.0426s

payoff would be 0.0511 as above. The value of swap to the seller would therefore be $0.0617 - 0.0511$, or 0.0106 times the principal. Similarly the mark-to-market value of the swap to the buyer of protection would be -0.0106 times the principal.

Estimating Default Probabilities

The default probabilities used to value a CDS should be risk-neutral default probabilities, not real-world default probabilities (see Section 23.5 for a discussion of the difference between the two). Risk-neutral default probabilities can be estimated from bond prices or asset swaps as explained in Chapter 23. An alternative is to imply them from CDS quotes. The latter approach is similar to the practice in options markets of implying volatilities from the prices of actively traded options.

Suppose we change the example in Tables 24.2, 24.3 and 24.4 so that we do not know the default probabilities. Instead we know that the mid-market CDS spread for a newly issued 5-year CDS is 100 basis points per year. We can reverse-engineer our calculations (using Excel in conjunction with Solver) to conclude that the implied default probability per year (conditional on no earlier default) is 1.61% per year.⁶

Binary Credit Default Swaps

A binary credit default swap is structured similarly to a regular credit default swap except that the payoff is a fixed dollar amount. Suppose that, in the example we considered in Tables 24.1 to 24.4, the payoff is \$1 instead of $1 - R$ dollars and the swap spread is s . Tables 24.1, 24.2 and 24.4 are the same, but Table 24.3 is replaced by Table 24.5. The CDS spread for a new binary CDS is given by $4.1130s = 0.0852$, so that the CDS spread, s , is 0.0207, or 207 basis points.

How Important is the Recovery Rate?

Whether we use CDS spreads or bond prices to estimate default probabilities we need an estimate of the recovery rate. However, provided that we use the same recovery rate

Table 24.5 Calculation of the present value of expected payoff from a binary credit default swap. Principal = \$1.

<i>Time (years)</i>	<i>Probability of default</i>	<i>Expected payoff (\$)</i>	<i>Discount factor</i>	<i>PV of expected payoff (\$)</i>
0.5	0.0200	0.0200	0.9753	0.0195
1.5	0.0196	0.0196	0.9277	0.0182
2.5	0.0192	0.0192	0.8825	0.0170
3.5	0.0188	0.0188	0.8395	0.0158
4.5	0.0184	0.0184	0.7985	0.0147
<i>Total</i>				0.0852

⁶ The DerivaGem worksheet gives a continuously compounded hazard rate of 1.626%. This is equivalent to 1.61% with annual compounding. If spreads for CDS swaps with different maturities are available, DerivaGem calculates a step function for the hazard rate.

for (a) estimating risk-neutral default probabilities and (b) valuing a CDS, the value of the CDS (or the estimate of the CDS spread) is not very sensitive to the recovery rate. This is because the implied probabilities of default are approximately proportional to $1/(1 - R)$ and the payoffs from a CDS are proportional to $1 - R$.

This argument does not apply to the valuation of binary CDS. Implied probabilities of default are still approximately proportional to $1/(1 - R)$. However, for a binary CDS, the payoffs from the CDS are independent of R . If we have a CDS spread for both a plain vanilla CDS and a binary CDS, we can estimate both the recovery rate and the default probability (see Problem 24.25).

The Future of the CDS Market

The credit default swap market survived the credit crunch of 2007 reasonably well. It is true that it has come under a great deal of regulatory scrutiny and CDSs are being moved to clearing houses. But their importance is unlikely to decline. They are important tools for managing credit risk. A financial institution can reduce its credit exposure to particular companies by buying protection. It can also use CDSs to diversify credit risk. For example, if a financial institution has too much credit exposure to a particular business sector, it can buy protection against defaults by companies in the sector and at the same time sell protection against default by companies in other unrelated sectors.

Some market participants think the CDS market will eventually be as big as the interest rate swap market. Others are less optimistic. There is a potential asymmetric information problem in the CDS market that is not present in other over-the-counter derivatives markets (see Business Snapshot 24.3).

24.3 CREDIT INDICES

Participants in credit markets have developed indices to track credit default swap spreads. In 2004 there were agreements between different producers of indices that led to some consolidation. Two important standard portfolios used by index providers are:

1. CDX NA IG, a portfolio of 125 investment grade companies in North America
2. iTraxx Europe, a portfolio of 125 investment grade names in Europe

These portfolios are updated on March 20 and September 20 each year. Companies that are no longer investment grade are dropped from the portfolios and new investment grade companies are added.⁷

Suppose that the 5-year CDX NA IG index is quoted by a market maker as bid 65 basis points, offer 66 basis points. (This is referred to as the index spread.) Roughly speaking, this means that a trader can buy CDS protection on all 125 companies in the index for 66 basis points per company. Suppose a trader wants \$800,000 of protection on each company. The total cost is $0.0066 \times 800,000 \times 125$, or \$660,000 per year. The trader can similarly sell \$800,000 of protection on each of the 125 companies for a total of \$650,000 per annum. When a company defaults, the protection buyer receives the

⁷ On September 20, 2010, the Series 14 iTraxx Europe portfolio and the Series 15 CDX NA IG portfolio were defined. The series numbers indicate that, by the end of September 2010, the iTraxx Europe portfolio had been updated 13 times and the CDX NA IG portfolio had been updated 14 times.

Business Snapshot 24.3 Is the CDS Market a Fair Game?

There is one important difference between credit default swaps and the other over-the-counter derivatives that we have considered in this book. The other over-the-counter derivatives depend on interest rates, exchange rates, equity indices, commodity prices, and so on. There is no reason to assume that any one market participant has better information than any other market participant about these variables.

Credit default swaps spreads depend on the probability that a particular company will default during a particular period of time. Arguably some market participants have more information to estimate this probability than others. A financial institution that works closely with a particular company by providing advice, making loans, and handling new issues of securities is likely to have more information about the creditworthiness of the company than another financial institution that has no dealings with the company. Economists refer to this as an *asymmetric information* problem.

Whether asymmetric information will curtail the expansion of the credit default swap market remains to be seen. Financial institutions emphasize that the decision to buy protection against the risk of default by a company is normally made by a risk manager and is not based on any special information that may exist elsewhere in the financial institution about the company.

usual CDS payoff and the annual payment is reduced by $660,000/125 = \$5,280$. There is an active market in buying and selling CDS index protection for maturities of 3, 5, 7, and 10 years. The maturities for these types of contracts on the index are usually December 20 and June 20. (This means that a “5-year” contract actually lasts between $4\frac{3}{4}$ and $5\frac{1}{4}$ years.) Roughly speaking, the index is the average of the CDS spreads on the companies in the underlying portfolio.⁸

24.4 THE USE OF FIXED COUPONS

The precise way in which CDS and CDS index transactions work is a little more complicated than has been described up to now. For each underlying and each maturity, a coupon and a recovery rate are specified. A price is calculated from the quoted spread using the following procedure:

1. Assume four payments per year, made in arrears.
2. Imply a hazard rate (default intensity) from the quoted spread. This involves calculations similar to those in Section 24.2. An iterative search is used to determine the hazard rate that leads to the quoted spread.

⁸ More precisely, the index is slightly lower than the average of the credit default swap spreads for the companies in the portfolio. To understand the reason for this consider a portfolio consisting of two companies, one with a spread of 1,000 basis points and the other with a spread of 10 basis points. To buy protection on the companies would cost slightly less than 505 basis points per company. This is because the 1,000 basis points is not expected to be paid for as long as the 10 basis points and should therefore carry less weight. Another complication for CDX NA IG, but not iTraxx Europe, is that the definition of default applicable to the index includes restructuring, whereas the definition for CDS contracts on the underlying companies often does not.

3. Calculate a duration D for the CDS payments. This is the number that the spread is multiplied by to get the present value of the spread payments. (In the example in Section 24.2, it is 4.1130.)
4. The price P is given by $P = 100 - 100 \times D \times (S - C)$, where S is the spread and C is the coupon expressed in decimal form.

When a trader buys protection the trader pays $100 - P$ per \$100 of the total remaining notional and the seller of protection receives this amount. (If $100 - P$ is negative, the buyer of protection receives money and the seller of protection pays money.) The buyer of protection then pays the coupon times the remaining notional on each payment date. (On a CDS, the remaining notional is the original notional until default and zero thereafter. For a CDS index, the remaining notional is the number of names in the index that have not yet defaulted multiplied by the principal per name.) The payoff when there is a default is calculated in the usual way. This arrangement facilitates trading because the regular quarterly payments made by the buyer of protection are independent of the spread at the time the buyer enters into the contract.

Example 24.1

Suppose that the iTraxx Europe index quote is 34 basis points and the coupon is 40 basis points for a contract lasting exactly 5 years, with both quotes being expressed using a 30/360 day count. (This is the usual day count convention in CDS and CDS index markets.) The equivalent actual/actual quotes are 0.345% for the index and 0.406% for the coupon. Suppose that the yield curve is flat at 4% per year (actual/actual, continuously compounded). The specified recovery rate is 40%. With four payments per year in arrears, the implied hazard rate is 0.5717%. The duration is 4.447 years. The price is therefore

$$100 - 100 \times 4.447 \times (0.00345 - 0.00406) = 100.27$$

Consider a contract where protection is \$1 million per name. Initially, the seller of protection would pay the buyer $\$1,000,000 \times 125 \times 0.0027$. Thereafter, the buyer of protection would make quarterly payments in arrears at an annual rate of $\$1,000,000 \times 0.00406 \times n$, where n is the number of companies that have not defaulted. When a company defaults, the payoff is calculated in the usual way and there is an accrual payment from the buyer to the seller calculated at the rate of 0.406% per year on \$1 million.

24.5 CDS FORWARDS AND OPTIONS

Once the CDS market was well established, it was natural for derivatives dealers to trade forwards and options on credit default swap spreads.⁹

A forward credit default swap is the obligation to buy or sell a particular credit default swap on a particular reference entity at a particular future time T . If the reference entity defaults before time T , the forward contract ceases to exist. Thus a bank could enter into a forward contract to sell 5-year protection on a company for

⁹ The valuation of these instruments is discussed in J.C. Hull and A. White, "The Valuation of Credit Default Swap Options," *Journal of Derivatives*, 10, 5 (Spring 2003): 40–50.

280 basis points starting in 1 year. If the company defaulted before the 1-year point, the forward contract would cease to exist.

A credit default swap option is an option to buy or sell a particular credit default swap on a particular reference entity at a particular future time T . For example, a trader could negotiate the right to buy 5-year protection on a company starting in 1 year for 280 basis points. This is a call option. If the 5-year CDS spread for the company in 1 year turns out to be more than 280 basis points, the option will be exercised; otherwise it will not be exercised. The cost of the option would be paid up front. Similarly an investor might negotiate the right to sell 5-year protection on a company for 280 basis points starting in 1 year. This is a put option. If the 5-year CDS spread for the company in 1 year turns out to be less than 280 basis points, the option will be exercised; otherwise it will not be exercised. Again the cost of the option would be paid up front. Like CDS forwards, CDS options are usually structured so that they cease to exist if the reference entity defaults before option maturity.

24.6 BASKET CREDIT DEFAULT SWAPS

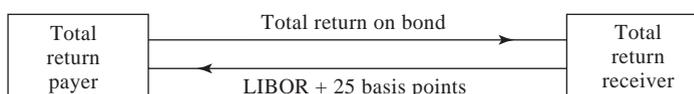
In what is referred to as a *basket credit default swap* there are a number of reference entities. An *add-up basket* CDS provides a payoff when any of the reference entities default. A *first-to-default* CDS provides a payoff only when the first default occurs. A *second-to-default* CDS provides a payoff only when the second default occurs. More generally, a *kth-to-default* CDS provides a payoff only when the k th default occurs. Payoffs are calculated in the same way as for a regular CDS. After the relevant default has occurred, there is a settlement. The swap then terminates and there are no further payments by either party.

24.7 TOTAL RETURN SWAPS

A *total return swap* is a type of credit derivative. It is an agreement to exchange the total return on a bond (or any portfolio of assets) for LIBOR plus a spread. The total return includes coupons, interest, and the gain or loss on the asset over the life of the swap.

An example of a total return swap is a 5-year agreement with a notional principal of \$100 million to exchange the total return on a corporate bond for LIBOR plus 25 basis points. This is illustrated in Figure 24.2. On coupon payment dates the payer pays the coupons earned on an investment of \$100 million in the bond. The receiver pays interest at a rate of LIBOR plus 25 basis points on a principal of \$100 million. (LIBOR is set on one coupon date and paid on the next as in a plain vanilla interest rate swap.) At the

Figure 24.2 Total return swap.



end of the life of the swap there is a payment reflecting the change in value of the bond. For example, if the bond increases in value by 10% over the life of the swap, the payer is required to pay \$10 million (= 10% of \$100 million) at the end of the 5 years. Similarly, if the bond decreases in value by 15%, the receiver is required to pay \$15 million at the end of the 5 years. If there is a default on the bond, the swap is usually terminated and the receiver makes a final payment equal to the excess of \$100 million over the market value of the bond.

If the notional principal is added to both sides at the end of the life of the swap, the total return swap can be characterized as follows. The payer pays the cash flows on an investment of \$100 million in the corporate bond. The receiver pays the cash flows on a \$100 million bond paying LIBOR plus 25 basis points. If the payer owns the corporate bond, the total return swap allows it to pass the credit risk on the bond to the receiver. If it does not own the bond, the total return swap allows it to take a short position in the bond.

Total return swaps are often used as a financing tool. One scenario that could lead to the swap in Figure 24.2 is as follows. The receiver wants financing to invest \$100 million in the reference bond. It approaches the payer (which is likely to be a financial institution) and agrees to the swap. The payer then invests \$100 million in the bond. This leaves the receiver in the same position as it would have been if it had borrowed money at LIBOR plus 25 basis points to buy the bond. The payer retains ownership of the bond for the life of the swap and faces less credit risk than it would have done if it had lent money to the receiver to finance the purchase of the bond, with the bond being used as collateral for the loan. If the receiver defaults the payer does not have the legal problem of trying to realize on the collateral. Total return swaps are similar to repos (see Section 4.1) in that they are structured to minimize credit risk when securities are being financed.

The spread over LIBOR received by the payer is compensation for bearing the risk that the receiver will default. The payer will lose money if the receiver defaults at a time when the reference bond's price has declined. The spread therefore depends on the credit quality of the receiver, the credit quality of the bond issuer, and the correlation between the two.

There are a number of variations on the standard deal we have described. Sometimes, instead of there being a cash payment for the change in value of the bond, there is physical settlement where the payer exchanges the underlying asset for the notional principal at the end of the life of the swap. Sometimes the change-in-value payments are made periodically rather than all at the end.

24.8 COLLATERAL DEBT OBLIGATIONS

We discussed asset-backed securities (ABSs) in Chapter 8. Figure 8.1 shows a simple structure. An ABS where the underlying assets are bonds is known as a *collateralized debt obligation*, or *CDO*. A waterfall similar to that indicated in Figure 8.2 is defined for the interest and principal payments on the bonds. The precise rules underlying the waterfall are complicated, but they are designed to ensure that if one tranche is more senior than another it is more likely to receive promised interest payments and repayments of principal.

Synthetic CDOs

When a CDO is created from a bond portfolio, as just described, the resulting structure is known as a *cash CDO*. In an important market development, it was recognized that a long position in a corporate bond has a similar risk to a short position in a CDS when the reference entity in the CDS is the company issuing the bond. This led an alternative structure known as a *synthetic CDO*, which has become very popular.

The originator of a synthetic CDO chooses a portfolio of companies and a maturity (e.g., 5 years) for the structure. It sells CDS protection on each company in the portfolio with the CDS maturities equaling the maturity of the structure. The synthetic CDO principal is the total of the notional principals underlying the CDSs. The originator has cash inflows equal to the the CDS spreads and cash outflows when companies in the portfolio default. Tranches are formed and the cash inflows and outflows are distributed to tranches. The rules for determining the cash inflows and outflows of tranches are more straightforward for a synthetic CDO than for a cash CDO. Suppose that there are only three tranches: equity, mezzanine, and senior. The rules might be as follows:

1. The equity tranche is responsible for the payouts on the CDSs until they reach 5% of the synthetic CDO principal. It earns a spread of 1000 basis points per year on the outstanding tranche principal.
2. The mezzanine tranche is responsible for payouts in excess of 5% up to a maximum of 20% of the synthetic CDO principal. It earns a spread of 100 basis points per year on the outstanding tranche principal.
3. The senior tranche is responsible for payouts in excess of 20%. It earns a spread of 10 basis points per year on the outstanding tranche principal.

To understand how the synthetic CDO would work, suppose that its principal is \$100 million. The equity, mezzanine, and senior tranche principals are \$5 million, \$15 million, and \$80 million, respectively. The tranches initially earn the specified spreads on these notional principals. Suppose that after 1 year defaults by companies in the portfolio lead to payouts of \$2 million on the CDSs. The equity tranche holders are responsible for these payouts. The equity tranche principal reduces to \$3 million and its spread (1,000 basis points) is then earned on \$3 million instead of \$5 million. If, later during the life of the CDO, there are further payouts of \$4 million on the CDSs, the cumulative of the payments required by the equity tranche is \$5 million, so that its outstanding principal becomes zero. The mezzanine tranche holders have to pay \$1 million. This reduces their outstanding principal to \$14 million.

Cash CDOs require an initial investment by the tranche holders (to finance the underlying bonds). By contrast, the holders of synthetic CDOs do not have to make an initial investment. They just have to agree to the way cash inflows and outflows will be calculated. In practice, they are almost invariably required to post the initial tranche principal as collateral. When the tranche becomes responsible for a payoff on a CDS, the money is taken out of the collateral. The balance in the collateral account earns interest at LIBOR.

Standard Portfolios and Single-Tranche Trading

In the synthetic CDO we have described, the tranche holders sell protection to the originator of the CDO, who in turn sells protection on CDSs to other market

Table 24.6 Mid-market quotes for 5-year tranches of iTraxx Europe. Quotes are in basis points except for the 0–3% tranche where the quote equals the percent of the tranche principal that must be paid up front in addition to 500 basis points per year. (Source: Creditex Group Inc.)

Date	Tranche					iTraxx index
	0–3%	3–6%	6–9%	9–12%	12–22%	
January 31, 2007	10.34%	41.59	11.95	5.60	2.00	23
January 31, 2008	30.98%	316.90	212.40	140.00	73.60	77
January 30, 2009	64.28%	1185.63	606.69	315.63	97.13	165

participants. An innovation in the market was the trading of a tranche without the underlying portfolio of short CDS positions being created. This is sometimes referred to as *single-tranche trading*. There are two parties to a trade: the buyer of protection on a tranche and the seller of protection on the tranche. The portfolio of short CDS positions is used as a reference point to define the cash flows between the two sides, but it is not created. The buyer of protection pays the tranche spread to the seller of protection, and the seller of protection pays amounts to the buyer that correspond to those losses on the reference portfolio of CDSs that the tranche is responsible for.

In Section 24.3, we discussed CDS indices such as CDX NA IG and iTraxx Europe. The market has used the portfolios underlying these indices to define standard synthetic CDO tranches. These trade very actively. The six standard tranches of CDX NA IG cover losses in the ranges 0–3%, 3–6%, 6–9%, 9–12%, 12–22%, and 22–100%. The six standard tranches of iTraxx Europe cover losses in the ranges 0–3%, 3–7%, 7–10%, 10–15%, 15–30%, and 30–100%.

Table 24.6 shows the quotes for 5-year iTraxx tranches at the end of January of three successive years. The index spread is the cost in basis points of buying protection on all the companies in the index, as described in Section 24.3. The quotes for all tranches except the 0–3% tranche is the cost in basis point per year of buying tranche protection. (As explained earlier, this is paid on a principal that declines as the tranche experiences losses.) In the case of the 0–3% (equity) tranche, the protection buyer makes an initial payment and then pays 500 basis points per year on the outstanding tranche principal. The quote is for the initial payment as a percentage of the initial tranche principal.

What a difference two years makes in the credit markets! Table 24.6 shows that the credit crunch led to a huge increase in credit spreads. The iTraxx index rose from 23 basis points in January 2007 to 165 basis points in January 2009. The individual tranche quotes have also shown huge increases. One reason for the changes is that the market's assessment of default probabilities for investment-grade corporations has increased. However, it is also the case that protection sellers were in many cases experiencing liquidity problems. They became more averse to risk and increased the risk premiums they required.

24.9 ROLE OF CORRELATION IN A BASKET CDS AND CDO

The cost of protection in a *k*th-to-default CDS or a tranche of a CDO is critically dependent on default correlation. Suppose that a basket of 100 reference entities is used

to define a 5-year k th-to-default CDS and that each reference entity has a risk-neutral probability of 2% of defaulting during the 5 years. When the default correlation between the reference entities is zero the binomial distribution shows that the probability of one or more defaults during the 5 years is 86.74% and the probability of 10 or more defaults is 0.0034%. A first-to-default CDS is therefore quite valuable whereas a tenth-to-default CDS is worth almost nothing.

As the default correlation increases the probability of one or more defaults declines and the probability of 10 or more defaults increases. In the limit where the default correlation between the reference entities is perfect the probability of one or more defaults equals the probability of ten or more defaults and is 2%. This is because in this extreme situation the reference entities are essentially the same. Either they all default (with probability 2%) or none of them default (with probability 98%).

The valuation of a tranche of a synthetic CDO is similarly dependent on default correlation. If the correlation is low, the junior equity tranche is very risky and the senior tranches are very safe. As the default correlation increases, the junior tranches become less risky and the senior tranches become more risky. In the limit where the default correlation is perfect and the recovery rate is zero, the tranches are equally risky.

24.10 VALUATION OF A SYNTHETIC CDO

Synthetic CDOs can be valued using the DerivaGem software. To explain the calculations, suppose that the payment dates on a synthetic CDO tranche are at times $\tau_1, \tau_2, \dots, \tau_m$ and $\tau_0 = 0$. Define E_j as the expected tranche principal at time τ_j and $v(\tau)$ as the present value of \$1 received at time τ . Suppose that the spread on a particular tranche (i.e., the number of basis points paid for protection) is s per year. This spread is paid on the remaining tranche principal. The present value of the expected regular spread payments on the CDO is therefore given by sA , where

$$A = \sum_{j=1}^m (\tau_j - \tau_{j-1}) E_j v(\tau_j) \quad (24.1)$$

The expected loss between times τ_{j-1} and τ_j is $E_{j-1} - E_j$. Assume that the loss occurs at the midpoint of the time interval (i.e., at time $0.5\tau_{j-1} + 0.5\tau_j$). The present value of the expected payoffs on the CDO tranche is

$$C = \sum_{j=1}^m (E_{j-1} - E_j) v(0.5\tau_{j-1} + 0.5\tau_j) \quad (24.2)$$

The accrual payment due on the losses is given by sB , where

$$B = \sum_{j=1}^m 0.5(\tau_j - \tau_{j-1})(E_{j-1} - E_j) v(0.5\tau_{j-1} + 0.5\tau_j) \quad (24.3)$$

The value of the tranche to the protection buyer is $C - sA - sB$. The breakeven spread on the tranche occurs when the present value of the payments equals the present value of the payoffs or

$$C = sA + sB$$

The breakeven spread is therefore

$$s = \frac{C}{A + B} \quad (24.4)$$

Equations (24.1) to (24.3) show the key role played by the expected tranche principal in calculating the breakeven spread for a tranche. If we know the expected principal for a tranche on all payment dates and we also know the zero-coupon yield curve, the breakeven tranche spread can be calculated from equations (24.1) to (24.4).

Using the Gaussian Copula Model of Time to Default

The one-factor Gaussian copula model of time to default was introduced in Section 22.9. This is the standard market model for valuing synthetic CDOs. All companies are assumed to have the same probability $Q(t)$ of defaulting by time t . Equation (23.12) converts this unconditional probability of default by time t to the probability of default by time t conditional on the factor F :

$$Q(t | F) = N\left(\frac{N^{-1}[Q(t)] - \sqrt{\rho} F}{\sqrt{1 - \rho}}\right) \quad (24.5)$$

Here ρ is the copula correlation, assumed to be the same for any pair of companies.

In the calculation of $Q(t)$, it is usually assumed that the hazard rate for a company is constant and consistent with the index spread. The hazard rate that is assumed can be calculated by using the CDS valuation approach in Section 24.2 and searching for the hazard rate that gives the index spread. Suppose that this hazard rate is λ . Then, from equation (23.1),

$$Q(t) = 1 - e^{-\lambda t} \quad (24.6)$$

From the properties of the binomial distribution, the standard market model gives the probability of exactly k defaults by time t , conditional on F , as

$$P(k, t | F) = \frac{n!}{(n - k)! k!} Q(t | F)^k [1 - Q(t | F)]^{n - k} \quad (24.7)$$

Suppose that the tranche under consideration covers losses on the portfolio between α_L and α_H . The parameter α_L is known as the *attachment point* and the parameter α_H is known as the *detachment point*. Define

$$n_L = \frac{\alpha_L n}{1 - R} \quad \text{and} \quad n_H = \frac{\alpha_H n}{1 - R}$$

where R is the recovery rate. Also, define $m(x)$ as the smallest integer greater than x . Without loss of generality, we assume that the initial tranche principal is 1. The tranche principal stays 1 while the number of defaults, k , is less than $m(n_L)$. It is zero when the number of defaults is greater than or equal to $m(n_H)$. Otherwise, the tranche principal is

$$\frac{\alpha_H - k(1 - R)/n}{\alpha_H - \alpha_L}$$

Define $E_j(F)$ as the expected tranche principal at time τ_j conditional on the value of the factor F . It follows that

$$E_j(F) = \sum_{k=0}^{m(n_L)-1} P(k, \tau_j | F) + \sum_{k=m(n_L)}^{m(n_H)-1} P(k, \tau_j | F) \frac{\alpha_H - k(1-R)/n}{\alpha_H - \alpha_L} \quad (24.8)$$

Define $A(F)$, $B(F)$, and $C(F)$ as the values of A , B , and C conditional on F . Similarly to equations (24.1) to (24.3),

$$A(F) = \sum_{j=1}^m (\tau_j - \tau_{j-1}) E_j(F) v(\tau_j) \quad (24.9)$$

$$B(F) = \sum_{j=1}^m 0.5(\tau_j - \tau_{j-1}) (E_{j-1}(F) - E_j(F)) v(0.5\tau_{j-1} + 0.5\tau_j) \quad (24.10)$$

$$C(F) = \sum_{j=1}^m (E_{j-1}(F) - E_j(F)) v(0.5\tau_{j-1} + 0.5\tau_j) \quad (24.11)$$

The variable F has a standard normal distribution. To calculate the unconditional values of A , B , and C , it is necessary to integrate $A(F)$, $B(F)$, and $C(F)$ over a standard normal distribution. Once the unconditional values have been calculated, the breakeven spread on the tranche can be calculated as $C/(A+B)$.¹⁰

The integration is best accomplished with a procedure known as Gaussian quadrature. It involves the following approximation:

$$\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-F^2/2} g(F) dF \approx \sum_{k=1}^{k=M} w_k g(F_k) \quad (24.12)$$

As M increases, accuracy increases. The values of w_k and F_k for different values of M are given on the author's website.¹¹ The value of M is twice the "number of integration points" variable in DerivaGem. Setting the number of integration points equal to 20 usually gives good results.

Example 24.2

Consider the mezzanine tranche of iTraxx Europe (5-year maturity) when the copula correlation is 0.15 and the recovery rate is 40%. In this case, $\alpha_L = 0.03$, $\alpha_H = 0.06$, $n = 125$, $n_L = 6.25$, and $n_H = 12.5$. We suppose that the term structure of interest rates is flat at 3.5%, payments are made quarterly, and the CDS spread on the index is 50 basis points. A calculation similar to that in Section 24.2 shows that the constant hazard rate corresponding to the CDS spread is 0.83% (with continuous compounding). An extract from the remaining calculations is shown in Table 24.7. A value of $M = 60$ is used in equation (24.12). The factor values, F_k , and their weights, w_k , are shown in first segment of the table. The expected tranche principals on payment dates conditional on the factor values are calculated from

¹⁰ In the case of the equity tranche, the quote is the upfront payment that must be made in addition to 500 basis points per year. The breakeven upfront payment is $C - 0.05(A+B)$.

¹¹ The parameters w_k and F_k are calculated from the roots of Hermite polynomials. For more information on Gaussian quadrature, see Technical Note 21 at www.rotman.utoronto.ca/~hull/TechnicalNotes.

Table 24.7 Valuation of CDO in Example 24.2: principal = 1; payments are per unit of spread.

Weights and values for factors

w_k	...	0.1579	0.1579	0.1342	0.0969	...
F_k	...	0.2020	-0.2020	-0.6060	-1.0104	...

Expected principal, $E_j(F_k)$

<i>Time</i>						
$j = 1$...	1.0000	1.0000	1.0000	1.0000	...
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
$j = 19$...	0.9953	0.9687	0.8636	0.6134	...
$j = 20$...	0.9936	0.9600	0.8364	0.5648	...

PV expected payment, $A(F_k)$

$j = 1$...	0.2478	0.2478	0.2478	0.2478	...
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
$j = 19$...	0.2107	0.2051	0.1828	0.1299	...
$j = 20$...	0.2085	0.2015	0.1755	0.1185	...
<i>Total</i>	...	4.5624	4.5345	4.4080	4.0361	...

PV expected accrual payment, $B(F_k)$

$j = 1$...	0.0000	0.0000	0.0000	0.0000	...
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
$j = 19$...	0.0001	0.0008	0.0026	0.0051	...
$j = 20$...	0.0002	0.0009	0.0029	0.0051	...
<i>Total</i>	...	0.0007	0.0043	0.0178	0.0478	...

PV expected payoff, $C(F_k)$

$j = 1$...	0.0000	0.0000	0.0000	0.0000	...
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
$j = 19$...	0.0011	0.0062	0.0211	0.0412	...
$j = 20$...	0.0014	0.0074	0.0230	0.0410	...
<i>Total</i>	...	0.0055	0.0346	0.1423	0.3823	...

equations (24.5) to (24.8) and shown in the second segment of the table. The values of A , B , and C conditional on the factor values are calculated in the last three segments of the table using equations (24.9) to (24.11). The unconditional values of A , B , and C are calculated by integrating $A(F)$, $B(F)$, and $C(F)$ over the probability distribution of F . This is done by setting $g(F)$ equal in turn to $A(F)$, $B(F)$, and $C(F)$ in equation (24.12). The result is

$$A = 4.2846, \quad B = 0.0187, \quad C = 0.1496$$

The breakeven tranche spread is $0.1496/(4.2846 + 0.0187) = 0.0348$, or 348 basis points.

This result can be obtained from DerivaGem. The CDS worksheet is used to convert the 50-basis-point spread to a hazard rate of 0.83%. The CDO worksheet is then used with this hazard rate and 30 integration points.

Valuation of k th-to-Default CDS

A k th-to-default CDS (see Section 24.5) can also be valued using the standard market model by conditioning on the factor F . The conditional probability that the k th default happens between times τ_{j-1} and τ_j is the probability that there are k or more defaults by time τ_j minus the probability that there are k or more defaults by time τ_{j-1} . This can be calculated from equations (24.5) to (24.7) as

$$\sum_{q=k}^n P(q, \tau_j | F) - \sum_{q=k}^n P(q, \tau_{j-1} | F)$$

Defaults between time τ_{j-1} and τ_j can be assumed to happen at time $0.5\tau_{j-1} + 0.5\tau_j$. This allows the present value of payments and of payoffs, conditional on F , to be calculated in the same way as for regular CDS payoffs (see Section 24.2). By integrating over F , the unconditional present values of payments and payoffs can be calculated.

Example 24.3

Consider a portfolio consisting of 10 bonds each with the default probabilities in Table 24.1 and suppose we are interested in valuing a third-to-default CDS where payments are made annually in arrears. Assume that the copula correlation is 0.3, the recovery rate is 40%, and all risk-free rates are 5%. As in Table 24.7, we consider $M = 60$ different factor values. The unconditional cumulative probability of each bond defaulting by years 1, 2, 3, 4, 5 is 0.0200, 0.0396, 0.0588, 0.0776, 0.0961, respectively. Equation (24.5) shows that, conditional on $F = -1.0104$, these default probabilities are 0.0365, 0.0754, 0.1134, 0.1498, 0.1848, respectively. From the binomial distribution, the conditional probability of three or more defaults by times 1, 2, 3, 4, 5 years is 0.0048, 0.0344, 0.0950, 0.1794, 0.2767, respectively. The conditional probability of the third default happening during years 1, 2, 3, 4, 5 is therefore 0.0048, 0.296, 0.0606, 0.0844, 0.0974, respectively. An analysis similar to that in Section 24.2 shows that the present values of regular payments, accrual payments, and payoffs conditional on $F = -1.0104$ are $3.8344s$, $0.1171s$, and 0.1405 , where s is the spread. Similar calculations are carried out for the other 59 factor values and equation (24.12) is used to integrate over F . The unconditional present values of payoffs, regular payments, and accrual payments are 0.0637, $4.0543s$, and $0.0531s$. The breakeven CDS spread is therefore $0.0637/(4.0543 + 0.0531) = 0.0155$, or 155 basis points.

Implied Correlation

In the standard market model, the recovery rate R is usually assumed to be 40%. This leaves the copula correlation ρ as the only unknown parameter. This makes the model similar to Black–Scholes–Merton, where there is only one unknown parameter, the volatility. Market participants like to imply a correlation from the market quotes for tranches in the same way that they imply a volatility from the market prices of options.

Suppose that the values of $\{\alpha_L, \alpha_H\}$ for successively more senior tranches are $\{\alpha_0, \alpha_1\}$, $\{\alpha_1, \alpha_2\}$, $\{\alpha_2, \alpha_3\}$, \dots , with $\alpha_0 = 0$. (For example, in the case of iTraxx Europe, $\alpha_0 = 0$, $\alpha_1 = 0.03$, $\alpha_2 = 0.06$, $\alpha_3 = 0.09$, $\alpha_4 = 0.12$, $\alpha_5 = 0.22$, $\alpha_6 = 1.00$.) There are two alternative implied correlations measures. One is *compound correlation*. For a tranche $\{\alpha_{q-1}, \alpha_q\}$, this is the value of the correlation, ρ , that leads to the spread

calculated from the model being the same as the spread in the market. It is found using an iterative search. The other is *base correlation*. For a particular value of α_q ($q \geq 1$), this is the value of ρ that leads to the $\{0, \alpha_q\}$ tranche being priced consistently with the market. It is obtained using the following steps:

1. Calculate the compound correlation for each tranche.
2. Use the compound correlation to calculate the present value of the expected loss on each tranche during the life of the CDO as a percent of the initial tranche principal. This is the variable we have defined as C above. Suppose that the value of C for the α_{q-1} to α_q tranche is C_q .
3. Calculate the present value of the expected loss on the $\{0, \alpha_q\}$ tranche as a percent of the total principal of the underlying portfolio. This is $\sum_{p=1}^q C_p(\alpha_p - \alpha_{p-1})$.
4. The C -value for the $\{0, \alpha_q\}$ tranche is the value calculated in Step 3 divided by α_q . The base correlation is the value of the correlation parameter, ρ , that is consistent with this C -value. It is found using an iterative search.

The present value of the loss as a percent of underlying portfolio that would be calculated in Step 3 for the iTraxx Europe quotes for January 31, 2007, given in Table 24.6 are shown in Figure 24.3. The implied correlations for these quotes are shown in Table 24.8. The calculations were carried out using DerivaGem assuming that the term structure of interest rates is flat at 3% and the recovery rate is 40%. The CDSs worksheet shows that the 23-basis-point spread implies a hazard rate of 0.382%. The implied correlations are calculated using the CDOs worksheet. The values underlying Figure 24.3 can also be calculated with this worksheet using the expression in Step 3 above.

The correlation patterns in Table 24.8 are typical of those usually observed. The compound correlations exhibit a “correlation smile”. As the tranche becomes more

Figure 24.3 Present value of expected loss on 0 to $X\%$ tranche as a percent of total underlying principal for iTraxx Europe on January 31, 2007.

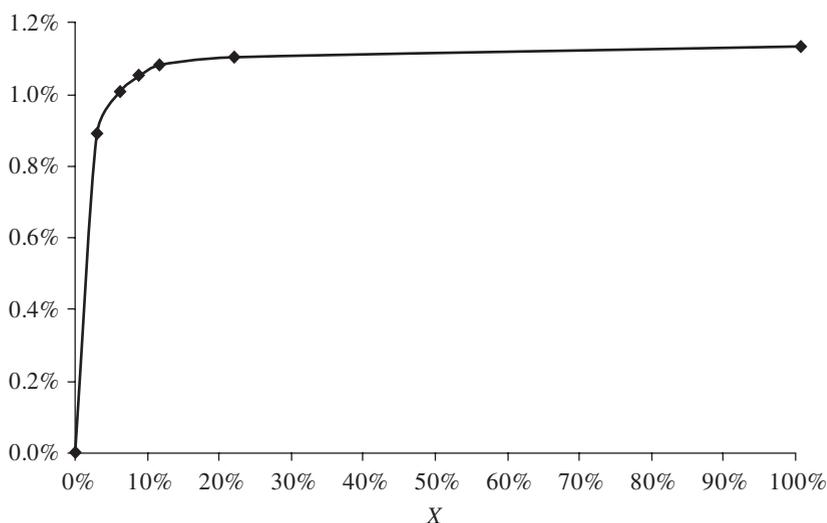


Table 24.8 Implied correlations for 5-year iTraxx Europe tranches on January 31, 2007.

Compound correlations					
Tranche	0–3%	3–6%	6–9%	9–12%	12–22%
Implied correlation	17.7%	7.8%	14.0%	18.2%	23.3%
Base correlations					
Tranche	0–3%	0–6%	0–9%	0–12%	0–22%
Implied correlation	17.7%	28.4%	36.5%	43.2%	60.5%

senior, the implied correlation first decreases and then increases. The base correlations exhibit a correlation skew where the implied correlation is an increasing function of the tranche detachment point.

If market prices were consistent with the one-factor Gaussian copula model, then the implied correlations (both compound and base) would be the same for all tranches. From the pronounced smiles and skews that are observed in practice, we can infer that market prices are not consistent with this model.

Valuing Nonstandard Tranches

We do not need a model to value the standard tranches of a standard portfolio such as iTraxx Europe because the spreads for these tranches can be observed in the market. Sometimes quotes need to be produced for nonstandard tranches of a standard portfolio. Suppose that you need a quote for the 4–8% iTraxx Europe tranche. One approach is to interpolate base correlations so as to estimate the base correlation for the 0–4% tranche and the 0–8% tranche. These two base correlations allow the present value of expected loss (as a percent of the underlying portfolio principal) to be estimated for these tranches. The present value of the expected loss for the 4–8% tranche (as a percent of the underlying principal) can be estimated as the difference between the present value of expected losses for the 0–8% and 0–4% tranches. This can be used to imply a compound correlation and a breakeven spread for the tranche.

It is now recognized that this is not the best way to proceed. A better approach is to calculate expected losses for each of the standard tranches and produce a chart such as Figure 24.3 showing the variation of expected loss for the 0– X % tranche with X . Values on this chart can be interpolated to give the expected loss for the 0–4% and the 0–8% tranches. The difference between these expected losses is a better estimate of the expected loss on the 4–8% tranche than that obtained from the base correlation approach.

It can be shown that for no arbitrage the expected losses in Figure 24.4 must increase at a decreasing rate. If base correlations are interpolated and then used to calculate expected losses, this no-arbitrage condition is often not satisfied. (The problem here is that the base correlation for the 0– X % tranche is a nonlinear function of the expected loss on the 0– X % tranche.) The direct approach of interpolating expected losses is therefore much better than the indirect approach of interpolating base correlations. What is more, it can be done so as to ensure that the no-arbitrage condition just mentioned is satisfied.

24.11 ALTERNATIVES TO THE STANDARD MARKET MODEL

This section outlines a number of alternatives to the one-factor Gaussian copula model that has become the market standard.

Heterogeneous Model

The standard market model is a homogeneous model in the sense that the time-to-default probability distributions are assumed to be the same for all companies and the copula correlations for any pair of companies are the same. The homogeneity assumption can be relaxed so that a more general model is used. However, this model is more complicated to implement because each company has a different probability of defaulting by any given time and $P(k, t | F)$ can no longer be calculated using the binomial formula in equation (24.7). It is necessary to use a numerical procedure such as that described in Andersen *et al.* (2003) and Hull and White (2004).¹²

Other Copulas

The one-factor Gaussian copula model is a particular model of the correlation between times to default. Many other one-factor copula models have been proposed. These include the Student t copula, the Clayton copula, Archimedean copula, and Marshall–Olkin copula. We can also create new one-factor copulas by assuming that F and the Z_i in equation (23.10) have nonnormal distributions with mean 0 and standard deviation 1. Hull and White show that a good fit to the market is obtained when F and the Z_i have Student t distributions with four degrees of freedom.¹³ They call this the *double t copula*.

Another approach is to increase the number of factors in the model. Unfortunately, the model is then much slower to run because it is necessary to integrate over several normal distributions instead of just one.

Random Factor Loadings

Andersen and Sidenius have suggested a model where the copula correlation ρ in equation (24.5) is a function of F .¹⁴

In general, ρ increases as F decreases. This means that in states of the world where the default rate is high (i.e., states of the world where F is low) the default correlation is also high. There is empirical evidence suggesting that this is the case.¹⁵ Andersen and Sidenius find that this model fits market quotes much better than the standard market model.

¹² See L. Andersen, J. Sidenius, and S. Basu, “All Your Hedges in One Basket,” *Risk*, November 2003; and J. C. Hull and A. White, “Valuation of a CDO and n th-to-Default Swap without Monte Carlo Simulation,” *Journal of Derivatives*, 12, 2 (Winter 2004), 8–23.

¹³ See J. C. Hull and A. White, “Valuation of a CDO and n th-to-Default Swap without Monte Carlo Simulation,” *Journal of Derivatives*, 12, 2 (Winter 2004), 8–23.

¹⁴ See L. Andersen and J. Sidenius, “Extension of the Gaussian Copula Model: Random Recovery and Random Factor Loadings,” *Journal of Credit Risk*, 1, 1 (Winter 2004), 29–70.

¹⁵ See, for example, A. Sevigny and O. Renault, “Default Correlation: Empirical Evidence,” Working Paper, Standard and Poors, 2002; S. R. Das, L. Freed, G. Geng, and N. Kapadia, “Correlated Default Risk,” *Journal of Fixed Income*, 16 (2006), 2, 7–32; J. C. Hull, M. Predescu, and A. White, “The Valuation of Correlation-Dependent Credit Derivatives Using a Structural Model,” *Journal of Credit Risk*, 6 (2010), 99–132; and A. Ang and J. Chen, “Asymmetric Correlation of Equity Portfolios,” *Journal of Financial Economics*, 63 (2002), 443–494.

The Implied Copula Model

Hull and White show how a copula can be implied from market quotes.¹⁶ The simplest version of the model assumes that a certain average hazard rate applies to all companies in a portfolio over the life of a CDO. That average hazard rate has a probability distribution that can be implied from the pricing of tranches. The calculation of the implied copula is similar in concept to the idea, discussed in Chapter 19, of calculating an implied probability distribution for a stock price from option prices.

Dynamic Models

The models discussed so far can be characterized as static models. In essence they model the average default environment over the life of the CDO. The model constructed for a 5-year CDO is different from the model constructed for a 7-year CDO, which is in turn different from the model constructed for a 10-year CDO. Dynamic models are different from static models in that they attempt to model the evolution of the loss on a portfolio through time. There are three different types of dynamic models:

1. *Structural Models*: These are similar to the models described in Section 23.6 except that the stochastic processes for the asset prices of many companies are modeled simultaneously. When the asset price for a company reaches a barrier, there is a default. The processes followed by the assets are correlated. The problem with these types of models is that they have to be implemented with Monte Carlo simulation and calibration is therefore difficult.
2. *Reduced Form Models*: In these models the hazard rates of companies are modeled. In order to build in a realistic amount of correlation, it is necessary to assume that there are jumps in the hazard rates.
3. *Top Down Models*: These are models where the total loss on a portfolio is modeled directly. The models do not consider what happens to individual companies.

SUMMARY

Credit derivatives enable banks and other financial institutions to actively manage their credit risks. They can be used to transfer credit risk from one company to another and to diversify credit risk by swapping one type of exposure for another.

The most common credit derivative is a credit default swap. This is a contract where one company buys insurance from another company against a third company (the reference entity) defaulting on its obligations. The payoff is usually the difference between the face value of a bond issued by the reference entity and its value immediately after a default. Credit default swaps can be analyzed by calculating the present value of the expected payments and the present value of the expected payoff in a risk-neutral world.

¹⁶ See J. C. Hull and A. White, "Valuing Credit Derivatives Using an Implied Copula Approach," *Journal of Derivatives*, 14 (2006), 8–28; and J. C. Hull and A. White, "An Improved Implied Copula Model and its Application to the Valuation of Bespoke CDO Tranches," *Journal of Investment Management*, 8, 3 (2010), 11–31.

A forward credit default swap is an obligation to enter into a particular credit default swap on a particular date. A credit default swap option is the right to enter into a particular credit default swap on a particular date. Both instruments cease to exist if the reference entity defaults before the date. A k th-to-default CDS is defined as a CDS that pays off when the k th default occurs in a portfolio of companies.

A total return swap is an instrument where the total return on a portfolio of credit-sensitive assets is exchanged for LIBOR plus a spread. Total return swaps are often used as financing vehicles. A company wanting to purchase a portfolio of assets will approach a financial institution to buy the assets on its behalf. The financial institution then enters into a total return swap with the company where it pays the return on the assets to the company and receives LIBOR plus a spread. The advantage of this type of arrangement is that the financial institution reduces its exposure to defaults by the company.

In a collateralized debt obligation a number of different securities are created from a portfolio of corporate bonds or commercial loans. There are rules for determining how credit losses are allocated. The result of the rules is that securities with both very high and very low credit ratings are created from the portfolio. A synthetic collateralized debt obligation creates a similar set of securities from credit default swaps. The standard market model for pricing both a k th-to-default CDS and tranches of a synthetic CDO is the one-factor Gaussian copula model for time to default.

FURTHER READING

- Andersen, L., and J. Sidenius, "Extensions to the Gaussian Copula: Random Recovery and Random Factor Loadings," *Journal of Credit Risk*, 1, No. 1 (Winter 2004): 29–70.
- Andersen, L., J. Sidenius, and S. Basu, "All Your Hedges in One Basket," *Risk*, November 2003.
- Das, S., *Credit Derivatives: Trading & Management of Credit & Default Risk*, 3rd edn. New York: Wiley, 2005.
- Hull, J. C., and A. White, "Valuation of a CDO and n th to Default Swap without Monte Carlo Simulation," *Journal of Derivatives*, 12, No. 2 (Winter 2004): 8–23.
- Hull, J. C., and A. White, "Valuing Credit Derivatives Using an Implied Copula Approach," *Journal of Derivatives*, 14, 2 (Winter 2006), 8–28.
- Hull, J. C., and A. White, "An Improved Implied Copula Model and its Application to the Valuation of Bespoke CDO Tranches," *Journal of Investment Management*, 8, 3 (2010), 11–31.
- Laurent, J.-P., and J. Gregory, "Basket Default Swaps, CDOs and Factor Copulas," *Journal of Risk*, 7, 4 (2005), 8–23.
- Li, D. X., "On Default Correlation: A Copula Approach," *Journal of Fixed Income*, March 2000: 43–54.
- Schönbucher, P. J., *Credit Derivatives Pricing Models*. New York: Wiley, 2003.
- Tavakoli, J. M., *Credit Derivatives & Synthetic Structures: A Guide to Instruments and Applications*, 2nd edn. New York: Wiley, 1998.

Practice Questions (Answers in Solutions Manual)

- 24.1. Explain the difference between a regular credit default swap and a binary credit default swap.
- 24.2. A credit default swap requires a semiannual payment at the rate of 60 basis points per year. The principal is \$300 million and the credit default swap is settled in cash.

A default occurs after 4 years and 2 months, and the calculation agent estimates that the price of the cheapest deliverable bond is 40% of its face value shortly after the default. List the cash flows and their timing for the seller of the credit default swap.

- 24.3. Explain the two ways a credit default swap can be settled.
- 24.4. Explain how a cash CDO and a synthetic CDO are created.
- 24.5. Explain what a first-to-default credit default swap is. Does its value increase or decrease as the default correlation between the companies in the basket increases? Explain.
- 24.6. Explain the difference between risk-neutral and real-world default probabilities.
- 24.7. Explain why a total return swap can be useful as a financing tool.
- 24.8. Suppose that the risk-free zero curve is flat at 7% per annum with continuous compounding and that defaults can occur halfway through each year in a new 5-year credit default swap. Suppose that the recovery rate is 30% and the default probabilities each year conditional on no earlier default is 3%. Estimate the credit default swap spread. Assume payments are made annually.
- 24.9. What is the value of the swap in Problem 24.8 per dollar of notional principal to the protection buyer if the credit default swap spread is 150 basis points?
- 24.10. What is the credit default swap spread in Problem 24.8 if it is a binary CDS?
- 24.11. How does a 5-year n th-to-default credit default swap work? Consider a basket of 100 reference entities where each reference entity has a probability of defaulting in each year of 1%. As the default correlation between the reference entities increases what would you expect to happen to the value of the swap when (a) $n = 1$ and (b) $n = 25$. Explain your answer.
- 24.12. What is the formula relating the payoff on a CDS to the notional principal and the recovery rate?
- 24.13. Show that the spread for a new plain vanilla CDS should be $(1 - R)$ times the spread for a similar new binary CDS, where R is the recovery rate.
- 24.14. Verify that, if the CDS spread for the example in Tables 24.1 to 24.4 is 100 basis points, the probability of default in a year (conditional on no earlier default) must be 1.61%. How does the probability of default change when the recovery rate is 20% instead of 40%? Verify that your answer is consistent with the implied probability of default being approximately proportional to $1/(1 - R)$, where R is the recovery rate.
- 24.15. A company enters into a total return swap where it receives the return on a corporate bond paying a coupon of 5% and pays LIBOR. Explain the difference between this and a regular swap where 5% is exchanged for LIBOR.
- 24.16. Explain how forward contracts and options on credit default swaps are structured.
- 24.17. “The position of a buyer of a credit default swap is similar to the position of someone who is long a risk-free bond and short a corporate bond.” Explain this statement.
- 24.18. Why is there a potential asymmetric information problem in credit default swaps?
- 24.19. Does valuing a CDS using real-world default probabilities rather than risk-neutral default probabilities overstate or understate its value? Explain your answer.
- 24.20. What is the difference between a total return swap and an asset swap?
- 24.21. Suppose that in a one-factor Gaussian copula model the 5-year probability of default for each of 125 names is 3% and the pairwise copula correlation is 0.2. Calculate, for factor

- values of -2 , -1 , 0 , 1 , and 2 : (a) the default probability conditional on the factor value and (b) the probability of more than 10 defaults conditional on the factor value.
- 24.22. Explain the difference between base correlation and compound correlation.
- 24.23. In Example 24.2, what is the tranche spread for the 9% to 12% tranche?

Further Questions

- 24.24. Suppose that the risk-free zero curve is flat at 6% per annum with continuous compounding and that defaults can occur at times 0.25 years, 0.75 years, 1.25 years, and 1.75 years in a 2-year plain vanilla credit default swap with semiannual payments. Suppose that the recovery rate is 20% and the unconditional probabilities of default (as seen at time zero) are 1% at times 0.25 years and 0.75 years, and 1.5% at times 1.25 years and 1.75 years. What is the credit default swap spread? What would the credit default spread be if the instrument were a binary credit default swap?
- 24.25. Assume that the default probability for a company in a year, conditional on no earlier defaults is λ and the recovery rate is R . The risk-free interest rate is 5% per annum. Default always occurs halfway through a year. The spread for a 5-year plain vanilla CDS where payments are made annually is 120 basis points and the spread for a 5-year binary CDS where payments are made annually is 160 basis points. Estimate R and λ .
- 24.26. Explain how you would expect the returns offered on the various tranches in a synthetic CDO to change when the correlation between the bonds in the portfolio increases.
- 24.27. Suppose that:
- (a) The yield on a 5-year risk-free bond is 7%.
 - (b) The yield on a 5-year corporate bond issued by company X is 9.5%.
 - (c) A 5-year credit default swap providing insurance against company X defaulting costs 150 basis points per year.
- What arbitrage opportunity is there in this situation? What arbitrage opportunity would there be if the credit default spread were 300 basis points instead of 150 basis points? Give two reasons why arbitrage opportunities such as those you identify are less than perfect.
- 24.28. In Example 24.3, what is the spread for (a) a first-to-default CDS and (b) a second-to-default CDS?
- 24.29. In Example 24.2, what is the tranche spread for the 6% to 9% tranche?
- 24.30. The 1-, 2-, 3-, 4-, and 5-year CDS spreads are 100, 120, 135, 145, and 152 basis points, respectively. The risk-free rate is 3% for all maturities, the recovery rate is 35%, and payments are quarterly. Use DerivaGem to calculate the continuously compounded hazard rate each year. What is the probability of default in year 1? What is the probability of default in year 2?
- 24.31. Table 24.6 shows the five-year iTraxx index was 77 basis points on January 31, 2008. Assume the risk-free rate is 5% for all maturities, the recovery rate is 40%, and payments are quarterly. Assume also that the spread of 77 basis points applies to all maturities. Use the DerivaGem CDS worksheet to calculate a hazard rate consistent with the spread. Use this in the CDO worksheet with 10 integration points to imply base correlations for each tranche from the quotes for January 31, 2008.



25

CHAPTER

Exotic Options

Derivatives such as European and American call and put options are what are termed *plain vanilla products*. They have standard well-defined properties and trade actively. Their prices or implied volatilities are quoted by exchanges or by inter-dealer brokers on a regular basis. One of the exciting aspects of the over-the-counter derivatives market is the number of nonstandard products that have been created by financial engineers. These products are termed *exotic options*, or simply *exotics*. Although they usually constitute a relatively small part of its portfolio, these exotics are important to a derivatives dealer because they are generally much more profitable than plain vanilla products.

Exotic products are developed for a number of reasons. Sometimes they meet a genuine hedging need in the market; sometimes there are tax, accounting, legal, or regulatory reasons why corporate treasurers, fund managers, and financial institutions find exotic products attractive; sometimes the products are designed to reflect a view on potential future movements in particular market variables; occasionally an exotic product is designed by a derivatives dealer to appear more attractive than it is to an unwary corporate treasurer or fund manager.

In this chapter, we describe some of the more commonly occurring exotic options and discuss their valuation. We assume that the asset provides a yield at rate q . As discussed in Chapters 16 and 17, for an option on a stock index q should be set equal to the dividend yield on the index, for an option on a currency it should be set equal to the foreign risk-free rate, and for an option on a futures contract it should be set equal to the domestic risk-free rate. Most of the options discussed in this chapter can be valued using the DerivaGem software.

25.1 PACKAGES

A *package* is a portfolio consisting of standard European calls, standard European puts, forward contracts, cash, and the underlying asset itself. We discussed a number of different types of packages in Chapter 11: bull spreads, bear spreads, butterfly spreads, calendar spreads, straddles, strangles, and so on.

Often a package is structured by traders so that it has zero cost initially. An example is a *range forward contract*.¹ This was discussed in Section 16.2. It consists of a long call and a short put or a short call and a long put. The call strike price is greater than the put strike price and the strike prices are chosen so that the value of the call equals the value of the put.

It is worth noting that any derivative can be converted into a zero-cost product by deferring payment until maturity. Consider a European call option. If c is the cost of the option when payment is made at time zero, then $A = ce^{rT}$ is the cost when payment is made at time T , the maturity of the option. The payoff is then $\max(S_T - K, 0) - A$ or $\max(S_T - K - A, -A)$. When the strike price, K , equals the forward price, other names for a deferred payment option are break forward, Boston option, forward with optional exit, and cancelable forward.

25.2 NONSTANDARD AMERICAN OPTIONS

In a standard American option, exercise can take place at any time during the life of the option and the exercise price is always the same. The American options that are traded in the over-the-counter market sometimes have nonstandard features. For example:

1. Early exercise may be restricted to certain dates. The instrument is then known as a *Bermudan option*. (Bermuda is between Europe and America!)
2. Early exercise may be allowed during only part of the life of the option. For example, there may be an initial “lock out” period with no early exercise.
3. The strike price may change during the life of the option.

The warrants issued by corporations on their own stock often have some or all of these features. For example, in a 7-year warrant, exercise might be possible on particular dates during years 3 to 7, with the strike price being \$30 during years 3 and 4, \$32 during the next 2 years, and \$33 during the final year.

Nonstandard American options can usually be valued using a binomial tree. At each node, the test (if any) for early exercise is adjusted to reflect the terms of the option.

25.3 GAP OPTIONS

A gap call option is a European call options that pays off $S_T - K_1$ when $S_T > K_2$. The difference between a gap call option and a regular call option with a strike price of K_2 is that the payoff when $S_T > K_2$ is increased by $K_2 - K_1$. (This increase is positive or negative depending on whether $K_2 > K_1$ or $K_1 > K_2$.)

A gap call option can be valued by a small modification to the Black–Scholes–Merton formula. With our usual notation, the value is

$$S_0 e^{-qT} N(d_1) - K_1 e^{-rT} N(d_2) \quad (25.1)$$

¹ Other names used for a range forward contract are zero-cost collar, flexible forward, cylinder option, option fence, min–max, and forward band.

where

$$d_1 = \frac{\ln(S_0/K_2) + (r - q + \sigma^2/2)T}{\sigma\sqrt{T}}$$

$$d_2 = d_1 - \sigma\sqrt{T}$$

The price in this formula is greater than the price given by the Black–Scholes–Merton formula for a regular call option with strike price K_2 by

$$(K_2 - K_1)e^{-rT}N(d_2)$$

To understand this difference, note that the probability that the option will be exercised is $N(d_2)$ and, when it is exercised, the payoff to the holder of the gap option is greater than that to the holder of the regular option by $K_2 - K_1$.

For a gap put option, the payoff is $K_1 - S_T$ when $S_T < K_2$. The value of the option is

$$K_1e^{-rT}N(-d_2) - S_0e^{-qT}N(-d_1) \quad (25.2)$$

where d_1 and d_2 are defined as for equation (25.1).

Example 25.1

An asset is currently worth \$500,000. Over the next year, it is expected to have a volatility of 20%. The risk-free rate is 5%, and no income is expected. Suppose that an insurance company agrees to buy the asset for \$400,000 if its value has fallen below \$400,000 at the end of one year. The payout will be $400,000 - S_T$ whenever the value of the asset is less than \$400,000. The insurance company has provided a regular put option where the policyholder has the right to sell the asset to the insurance company for \$400,000 in one year. This can be valued using equation (14.21), with $S_0 = 500,000$, $K = 400,000$, $r = 0.05$, $\sigma = 0.2$, $T = 1$. The value is \$3,436.

Suppose next that the cost of transferring the asset is \$50,000 and this cost is borne by the policyholder. The option is then exercised only if the value of the asset is less than \$350,000. In this case, the cost to the insurance company is $K_1 - S_T$ when $S_T < K_2$, where $K_2 = 350,000$, $K_1 = 400,000$, and S_T is the price of the asset in one year. This is a gap put option. The value is given by equation (25.2), with $S_0 = 500,000$, $K_1 = 400,000$, $K_2 = 350,000$, $r = 0.05$, $q = 0$, $\sigma = 0.2$, $T = 1$. It is \$1,896. Recognizing the costs to the policyholder of making a claim reduces the cost of the policy to the insurance company by about 45% in this case.

25.4 FORWARD START OPTIONS

Forward start options are options that will start at some time in the future. Sometimes employee stock options, which were discussed in Chapter 15, can be viewed as forward start options. This is because the company commits (implicitly or explicitly) to granting at-the-money options to employees in the future.

Consider a forward start at-the-money European call option that will start at time T_1 and mature at time T_2 . Suppose that the asset price is S_0 at time zero and S_1 at time T_1 .

To value the option, we note from the European option pricing formulas in Chapters 14 and 16 that the value of an at-the-money call option on an asset is proportional to the asset price. The value of the forward start option at time T_1 is therefore cS_1/S_0 , where c is the value at time zero of an at-the-money option that lasts for $T_2 - T_1$. Using risk-neutral valuation, the value of the forward start option at time zero is

$$e^{-rT_1} \hat{E} \left[c \frac{S_1}{S_0} \right]$$

where \hat{E} denotes the expected value in a risk-neutral world. Since c and S_0 are known and $\hat{E}[S_1] = S_0 e^{(r-q)T_1}$, the value of the forward start option is ce^{-qT_1} . For a non-dividend-paying stock, $q = 0$ and the value of the forward start option is exactly the same as the value of a regular at-the-money option with the same life as the forward start option.

25.5 CLIQUET OPTIONS

A cliquet option (which is also called a ratchet or strike reset option) is a series of call or put options with rules for determining the strike price. Suppose that the reset dates are at times $\tau, 2\tau, \dots, (n-1)\tau$, with $n\tau$ being the end of the cliquet's life. A simple structure would be as follows. The first option has a strike price K (which might equal the initial asset price) and lasts between times 0 and τ ; the second option provides a payoff at time 2τ with a strike price equal to the value of the asset at time τ ; the third option provides a payoff at time 3τ with a strike price equal to the value of the asset at time 2τ ; and so on. This is a regular option plus $n - 1$ forward start options. The latter can be valued as described in Section 25.2.

Some cliquet options are much more complicated than the one described here. For example, sometimes there are upper and lower limits on the total payoff over the whole period; sometimes cliquets terminate at the end of a period if the asset price is in a certain range. When analytic results are not available, Monte Carlo simulation can often be used for valuation.

25.6 COMPOUND OPTIONS

Compound options are options on options. There are four main types of compound options: a call on a call, a put on a call, a call on a put, and a put on a put. Compound options have two strike prices and two exercise dates. Consider, for example, a call on a call. On the first exercise date, T_1 , the holder of the compound option is entitled to pay the first strike price, K_1 , and receive a call option. The call option gives the holder the right to buy the underlying asset for the second strike price, K_2 , on the second exercise date, T_2 . The compound option will be exercised on the first exercise date only if the value of the option on that date is greater than the first strike price.

When the usual geometric Brownian motion assumption is made, European-style compound options can be valued analytically in terms of integrals of the bivariate normal distribution.² With our usual notation, the value at time zero of a European call

² See R. Geske, "The Valuation of Compound Options," *Journal of Financial Economics*, 7 (1979): 63–81; M. Rubinstein, "Double Trouble," *Risk*, December 1991/January 1992: 53–56.

option on a call option is

$$S_0 e^{-qT_2} M(a_1, b_1; \sqrt{T_1/T_2}) - K_2 e^{-rT_2} M(a_2, b_2; \sqrt{T_1/T_2}) - e^{-rT_1} K_1 N(a_2)$$

where

$$a_1 = \frac{\ln(S_0/S^*) + (r - q + \sigma^2/2)T_1}{\sigma\sqrt{T_1}}, \quad a_2 = a_1 - \sigma\sqrt{T_1}$$

$$b_1 = \frac{\ln(S_0/K_2) + (r - q + \sigma^2/2)T_2}{\sigma\sqrt{T_2}}, \quad b_2 = b_1 - \sigma\sqrt{T_2}$$

The function $M(a, b; \rho)$ is the cumulative bivariate normal distribution function that the first variable will be less than a and the second will be less than b when the coefficient of correlation between the two is ρ .³ The variable S^* is the asset price at time T_1 for which the option price at time T_1 equals K_1 . If the actual asset price is above S^* at time T_1 , the first option will be exercised; if it is not above S^* , the option expires worthless.

With similar notation, the value of a European put on a call is

$$K_2 e^{-rT_2} M(-a_2, b_2; -\sqrt{T_1/T_2}) - S_0 e^{-qT_2} M(-a_1, b_1; -\sqrt{T_1/T_2}) + e^{-rT_1} K_1 N(-a_2)$$

The value of a European call on a put is

$$K_2 e^{-rT_2} M(-a_2, -b_2; \sqrt{T_1/T_2}) - S_0 e^{-qT_2} M(-a_1, -b_1; \sqrt{T_1/T_2}) - e^{-rT_1} K_1 N(-a_2)$$

The value of a European put on a put is

$$S_0 e^{-qT_2} M(a_1, -b_1; -\sqrt{T_1/T_2}) - K_2 e^{-rT_2} M(a_2, -b_2; -\sqrt{T_1/T_2}) + e^{-rT_1} K_1 N(a_2)$$

25.7 CHOOSER OPTIONS

A *chooser* option (sometimes referred to as an *as you like it* option) has the feature that, after a specified period of time, the holder can choose whether the option is a call or a put. Suppose that the time when the choice is made is T_1 . The value of the chooser option at this time is

$$\max(c, p)$$

where c is the value of the call underlying the option and p is the value of the put underlying the option.

If the options underlying the chooser option are both European and have the same strike price, put–call parity can be used to provide a valuation formula. Suppose that S_1 is the asset price at time T_1 , K is the strike price, T_2 is the maturity of the options, and r is the risk-free interest rate. Put–call parity implies that

$$\begin{aligned} \max(c, p) &= \max(c, c + Ke^{-r(T_2-T_1)} - S_1 e^{-q(T_2-T_1)}) \\ &= c + e^{-q(T_2-T_1)} \max(0, Ke^{-(r-q)(T_2-T_1)} - S_1) \end{aligned}$$

³ See Technical Note 5 at www.rotman.utoronto.ca/~hull/TechnicalNotes for a numerical procedure for calculating M . A function for calculating M is also on the website.

This shows that the chooser option is a package consisting of:

1. A call option with strike price K and maturity T_2
2. $e^{-q(T_2-T_1)}$ put options with strike price $Ke^{-(r-q)(T_2-T_1)}$ and maturity T_1

As such, it can readily be valued.

More complex chooser options can be defined where the call and the put do not have the same strike price and time to maturity. They are then not packages and have features that are somewhat similar to compound options.

25.8 BARRIER OPTIONS

Barrier options are options where the payoff depends on whether the underlying asset's price reaches a certain level during a certain period of time.

A number of different types of barrier options regularly trade in the over-the-counter market. They are attractive to some market participants because they are less expensive than the corresponding regular options. These barrier options can be classified as either *knock-out options* or *knock-in options*. A knock-out option ceases to exist when the underlying asset price reaches a certain barrier; a knock-in option comes into existence only when the underlying asset price reaches a barrier.

Equations (16.4) and (16.5) show that the values at time zero of a regular call and put option are

$$\begin{aligned}c &= S_0 e^{-qT} N(d_1) - K e^{-rT} N(d_2) \\ p &= K e^{-rT} N(-d_2) - S_0 e^{-qT} N(-d_1)\end{aligned}$$

where

$$\begin{aligned}d_1 &= \frac{\ln(S_0/K) + (r - q + \sigma^2/2)T}{\sigma\sqrt{T}} \\ d_2 &= \frac{\ln(S_0/K) + (r - q - \sigma^2/2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}\end{aligned}$$

A *down-and-out call* is one type of knock-out option. It is a regular call option that ceases to exist if the asset price reaches a certain barrier level H . The barrier level is below the initial asset price. The corresponding knock-in option is a *down-and-in call*. This is a regular call that comes into existence only if the asset price reaches the barrier level.

If H is less than or equal to the strike price, K , the value of a down-and-in call at time zero is

$$c_{di} = S_0 e^{-qT} (H/S_0)^{2\lambda} N(y) - K e^{-rT} (H/S_0)^{2\lambda-2} N(y - \sigma\sqrt{T})$$

where

$$\begin{aligned}\lambda &= \frac{r - q + \sigma^2/2}{\sigma^2} \\ y &= \frac{\ln[H^2/(S_0 K)]}{\sigma\sqrt{T}} + \lambda\sigma\sqrt{T}\end{aligned}$$

Because the value of a regular call equals the value of a down-and-in call plus the value of a down-and-out call, the value of a down-and-out call is given by

$$c_{do} = c - c_{di}$$

If $H \geq K$, then

$$c_{do} = S_0 N(x_1) e^{-qT} - K e^{-rT} N(x_1 - \sigma\sqrt{T}) \\ - S_0 e^{-qT} (H/S_0)^{2\lambda} N(y_1) + K e^{-rT} (H/S_0)^{2\lambda-2} N(y_1 - \sigma\sqrt{T})$$

and

$$c_{di} = c - c_{do}$$

where

$$x_1 = \frac{\ln(S_0/H)}{\sigma\sqrt{T}} + \lambda\sigma\sqrt{T} \\ y_1 = \frac{\ln(H/S_0)}{\sigma\sqrt{T}} + \lambda\sigma\sqrt{T}$$

An *up-and-out call* is a regular call option that ceases to exist if the asset price reaches a barrier level, H , that is higher than the current asset price. An *up-and-in call* is a regular call option that comes into existence only if the barrier is reached. When H is less than or equal to K , the value of the up-and-out call, c_{uo} , is zero and the value of the up-and-in call, c_{ui} , is c . When H is greater than K ,

$$c_{ui} = S_0 N(x_1) e^{-qT} - K e^{-rT} N(x_1 - \sigma\sqrt{T}) - S_0 e^{-qT} (H/S_0)^{2\lambda} [N(-y) - N(-y_1)] \\ + K e^{-rT} (H/S_0)^{2\lambda-2} [N(-y + \sigma\sqrt{T}) - N(-y_1 + \sigma\sqrt{T})]$$

and

$$c_{uo} = c - c_{ui}$$

Put barrier options are defined similarly to call barrier options. An *up-and-out put* is a put option that ceases to exist when a barrier, H , that is greater than the current asset price is reached. An *up-and-in put* is a put that comes into existence only if the barrier is reached. When the barrier, H , is greater than or equal to the strike price, K , their prices are

$$p_{ui} = -S_0 e^{-qT} (H/S_0)^{2\lambda} N(-y) + K e^{-rT} (H/S_0)^{2\lambda-2} N(-y + \sigma\sqrt{T})$$

and

$$p_{uo} = p - p_{ui}$$

When H is less than or equal to K ,

$$p_{uo} = -S_0 N(-x_1) e^{-qT} + K e^{-rT} N(-x_1 + \sigma\sqrt{T}) \\ + S_0 e^{-qT} (H/S_0)^{2\lambda} N(-y_1) - K e^{-rT} (H/S_0)^{2\lambda-2} N(-y_1 + \sigma\sqrt{T})$$

and

$$p_{ui} = p - p_{uo}$$

A *down-and-out put* is a put option that ceases to exist when a barrier less than the current asset price is reached. A *down-and-in put* is a put option that comes into

existence only when the barrier is reached. When the barrier is greater than the strike price, $p_{do} = 0$ and $p_{di} = p$. When the barrier is less than the strike price,

$$p_{di} = -S_0 N(-x_1) e^{-qT} + K e^{-rT} N(-x_1 + \sigma\sqrt{T}) + S_0 e^{-qT} (H/S_0)^{2\lambda} [N(y) - N(y_1)] \\ - K e^{-rT} (H/S_0)^{2\lambda-2} [N(y - \sigma\sqrt{T}) - N(y_1 - \sigma\sqrt{T})]$$

and

$$p_{do} = p - p_{di}$$

All of these valuations make the usual assumption that the probability distribution for the asset price at a future time is lognormal. An important issue for barrier options is the frequency with which the asset price, S , is observed for purposes of determining whether the barrier has been reached. The analytic formulas given in this section assume that S is observed continuously and sometimes this is the case.⁴ Often, the terms of a contract state that S is observed periodically; for example, once a day at 12 noon. Broadie, Glasserman, and Kou provide a way of adjusting the formulas we have just given for the situation where the price of the underlying is observed discretely.⁵ The barrier level H is replaced by $He^{0.5826\sigma\sqrt{T/m}}$ for an up-and-in or up-and-out option and by $He^{-0.5826\sigma\sqrt{T/m}}$ for a down-and-in or down-and-out option, where m is the number of times the asset price is observed (so that T/m is the time interval between observations).

Barrier options often have quite different properties from regular options. For example, sometimes vega is negative. Consider an up-and-out call option when the asset price is close to the barrier level. As volatility increases, the probability that the barrier will be hit increases. As a result, a volatility increase can cause the price of the barrier option to decrease in these circumstances.

One disadvantage of the barrier options we have considered so far is that a “spike” in the asset price can cause the option to be knocked in or out. An alternative structure is a *Parisian option*, where the asset price has to be above or below the barrier for a period of time for the option to be knocked in or out. For example, a down-and-out Parisian put option with a strike price equal to 90% of the initial asset price and a barrier at 75% of the initial asset price might specify that the option is knocked out if the asset price is below the barrier for 50 days. The confirmation might specify that the 50 days are a “continuous period of 50 days” or “any 50 days during the option’s life.” Parisian options are more difficult to value than regular barrier options.⁶ Monte Carlo simulation and binomial trees can be used with the enhancements discussed in Section 26.5 and 26.6.

25.9 BINARY OPTIONS

Binary options are options with discontinuous payoffs. A simple example of a binary option is a *cash-or-nothing call*. This pays off nothing if the asset price ends up below the strike price at time T and pays a fixed amount, Q , if it ends up above the strike

⁴ One way to track whether a barrier has been reached from below (above) is to send a limit order to the exchange to sell (buy) the asset at the barrier price and see whether the order is filled.

⁵ M. Broadie, P. Glasserman, and S. G. Kou, “A Continuity Correction for Discrete Barrier Options,” *Mathematical Finance* 7, 4 (October 1997): 325–49.

⁶ See, for example, M. Chesney, J. Cornwall, M. Jeanblanc-Picque, G. Kentwell, and M. Yor, “Parisian pricing,” *Risk*, 10, 1 (1977), 77–79.

price. In a risk-neutral world, the probability of the asset price being above the strike price at the maturity of an option is, with our usual notation, $N(d_2)$. The value of a cash-or-nothing call is therefore $Qe^{-rT}N(d_2)$. A *cash-or-nothing put* is defined analogously to a cash-or-nothing call. It pays off Q if the asset price is below the strike price and nothing if it is above the strike price. The value of a cash-or-nothing put is $Qe^{-rT}N(-d_2)$.

Another type of binary option is an *asset-or-nothing call*. This pays off nothing if the underlying asset price ends up below the strike price and pays the asset price if it ends up above the strike price. With our usual notation, the value of an asset-or-nothing call is $S_0e^{-qT}N(d_1)$. An *asset-or-nothing put* pays off nothing if the underlying asset price ends up above the strike price and the asset price if it ends up below the strike price. The value of an asset-or-nothing put is $S_0e^{-qT}N(-d_1)$.

A regular European call option is equivalent to a long position in an asset-or-nothing call and a short position in a cash-or-nothing call where the cash payoff in the cash-or-nothing call equals the strike price. Similarly, a regular European put option is equivalent to a long position in a cash-or-nothing put and a short position in an asset-or-nothing put where the cash payoff on the cash-or-nothing put equals the strike price.

25.10 LOOKBACK OPTIONS

The payoffs from lookback options depend on the maximum or minimum asset price reached during the life of the option. The payoff from a *floating lookback call* is the amount that the final asset price exceeds the minimum asset price achieved during the life of the option. The payoff from a *floating lookback put* is the amount by which the maximum asset price achieved during the life of the option exceeds the final asset price.

Valuation formulas have been produced for floating lookbacks.⁷ The value of a floating lookback call at time zero is

$$c_{\bar{n}} = S_0e^{-qT}N(a_1) - S_0e^{-qT}\frac{\sigma^2}{2(r-q)}N(-a_1) - S_{\min}e^{-rT}\left[N(a_2) - \frac{\sigma^2}{2(r-q)}e^{Y_1}N(-a_3)\right]$$

where

$$\begin{aligned} a_1 &= \frac{\ln(S_0/S_{\min}) + (r - q + \sigma^2/2)T}{\sigma\sqrt{T}} \\ a_2 &= a_1 - \sigma\sqrt{T}, \\ a_3 &= \frac{\ln(S_0/S_{\min}) + (-r + q + \sigma^2/2)T}{\sigma\sqrt{T}} \\ Y_1 &= -\frac{2(r - q - \sigma^2/2)\ln(S_0/S_{\min})}{\sigma^2} \end{aligned}$$

and S_{\min} is the minimum asset price achieved to date. (If the lookback has just been originated, $S_{\min} = S_0$.) See Problem 25.23 for the $r = q$ case.

⁷ See B. Goldman, H. Sosin, and M. A. Gatto, "Path-Dependent Options: Buy at the Low, Sell at the High," *Journal of Finance*, 34 (December 1979): 1111-27.; M. Garman, "Recollection in Tranquility," *Risk*, March (1989): 16-19.

The value of a floating lookback put is

$$p_{\text{fl}} = S_{\text{max}} e^{-rT} \left[N(b_1) - \frac{\sigma^2}{2(r-q)} e^{Y_2} N(-b_3) \right] + S_0 e^{-qT} \frac{\sigma^2}{2(r-q)} N(-b_2) - S_0 e^{-qT} N(b_2)$$

where

$$b_1 = \frac{\ln(S_{\text{max}}/S_0) + (-r + q + \sigma^2/2)T}{\sigma\sqrt{T}}$$

$$b_2 = b_1 - \sigma\sqrt{T}$$

$$b_3 = \frac{\ln(S_{\text{max}}/S_0) + (r - q - \sigma^2/2)T}{\sigma\sqrt{T}}$$

$$Y_2 = \frac{2(r - q - \sigma^2/2) \ln(S_{\text{max}}/S_0)}{\sigma^2}$$

and S_{max} is the maximum asset price achieved to date. (If the lookback has just been originated, then $S_{\text{max}} = S_0$.)

A floating lookback call is a way that the holder can buy the underlying asset at the lowest price achieved during the life of the option. Similarly, a floating lookback put is a way that the holder can sell the underlying asset at the highest price achieved during the life of the option.

Example 25.2

Consider a newly issued floating lookback put on a non-dividend-paying stock where the stock price is 50, the stock price volatility is 40% per annum, the risk-free rate is 10% per annum, and the time to maturity is 3 months. In this case, $S_{\text{max}} = 50$, $S_0 = 50$, $r = 0.1$, $q = 0$, $\sigma = 0.4$, and $T = 0.25$, $b_1 = -0.025$, $b_2 = -0.225$, $b_3 = 0.025$, and $Y_2 = 0$, so that the value of the lookback put is 7.79. A newly issued floating lookback call on the same stock is worth 8.04.

In a fixed lookback option, a strike price is specified. For a *fixed lookback call option*, the payoff is the same as a regular European call option except that the final asset price is replaced by the maximum asset price achieved during the life of the option. For a *fixed lookback put option*, the payoff is the same as a regular European put option except that the the final asset price is replaced by the minimum asset price achieved during the life of the option. Define $S_{\text{max}}^* = \max(S_{\text{max}}, K)$, where as before S_{max} is the maximum asset price achieved to date and K is the strike price. Also, define p_{fl}^* as the value of a floating lookback put which lasts for the same period as the fixed lookback call when the actual maximum asset price so far, S_{max} , is replaced by S_{max}^* . A put-call parity type of argument shows that the value of the fixed lookback call option, c_{fix} is given by⁸

$$c_{\text{fix}} = p_{\text{fl}}^* + S_0 e^{-qT} - K e^{-rT}$$

Similarly, if $S_{\text{min}}^* = \min(S_{\text{min}}, K)$, then the value of a fixed lookback put option, p_{fix} , is given by

$$p_{\text{fix}} = c_{\text{fl}}^* + K e^{-rT} - S_0 e^{-qT}$$

⁸ The argument was proposed by H. Y. Wong and Y. K. Kwok, "Sub-replication and Replenishing Premium: Efficient Pricing of Multi-state Lookbacks," *Review of Derivatives Research*, 6 (2003), 83–106.

where c_{fl}^* is the value of a floating lookback call that lasts for the same period as the fixed lookback put when the actual minimum asset price so far, S_{min} , is replaced by S_{min}^* . This shows that the equations given above for floating lookbacks can be modified to price fixed lookbacks.

Lookbacks are appealing to investors, but very expensive when compared with regular options. As with barrier options, the value of a lookback option is liable to be sensitive to the frequency with which the asset price is observed for the purposes of computing the maximum or minimum. The formulas above assume that the asset price is observed continuously. Broadie, Glasserman, and Kou provide a way of adjusting the formulas we have just given for the situation where the asset price is observed discretely.⁹

25.11 SHOUT OPTIONS

A *shout option* is a European option where the holder can “shout” to the writer at one time during its life. At the end of the life of the option, the option holder receives either the usual payoff from a European option or the intrinsic value at the time of the shout, whichever is greater. Suppose the strike price is \$50 and the holder of a call shouts when the price of the underlying asset is \$60. If the final asset price is less than \$60, the holder receives a payoff of \$10. If it is greater than \$60, the holder receives the excess of the asset price over \$50.

A shout option has some of the same features as a lookback option, but is considerably less expensive. It can be valued by noting that if the holder shouts at a time τ when the asset price is S_τ the payoff from the option is

$$\max(0, S_T - S_\tau) + (S_\tau - K)$$

where, as usual, K is the strike price and S_T is the asset price at time T . The value at time τ if the holder shouts is therefore the present value of $S_\tau - K$ (received at time T) plus the value of a European option with strike price S_τ . The latter can be calculated using Black–Scholes–Merton formulas.

A shout option is valued by constructing a binomial or trinomial tree for the underlying asset in the usual way. Working back through the tree, the value of the option if the holder shouts and the value if the holder does not shout can be calculated at each node. The option’s price at the node is the greater of the two. The procedure for valuing a shout option is therefore similar to the procedure for valuing a regular American option.

25.12 ASIAN OPTIONS

Asian options are options where the payoff depends on the arithmetic average of the price of the underlying asset during the life of the option. The payoff from an *average price call* is $\max(0, S_{\text{ave}} - K)$ and that from an *average price put* is $\max(0, K - S_{\text{ave}})$, where S_{ave} is the average price of the underlying asset. Average price options are less

⁹ M. Broadie, P. Glasserman, and S.G. Kou, “Connecting Discrete and Continuous Path-Dependent Options,” *Finance and Stochastics*, 2 (1998): 1–28.

expensive than regular options and are arguably more appropriate than regular options for meeting some of the needs of corporate treasurers. Suppose that a US corporate treasurer expects to receive a cash flow of 100 million Australian dollars spread evenly over the next year from the company's Australian subsidiary. The treasurer is likely to be interested in an option that guarantees that the average exchange rate realized during the year is above some level. An average price put option can achieve this more effectively than regular put options.

Average price options can be valued using similar formulas to those used for regular options if it is assumed that S_{ave} is lognormal. As it happens, when the usual assumption is made for the process followed by the asset price, this is a reasonable assumption.¹⁰ A popular approach is to fit a lognormal distribution to the first two moments of S_{ave} and use Black's model.¹¹ Suppose that M_1 and M_2 are the first two moments of S_{ave} . The value of average price calls and puts are given by equations (17.9) and (17.10), with

$$F_0 = M_1 \quad (25.3)$$

and

$$\sigma^2 = \frac{1}{T} \ln \left(\frac{M_2}{M_1^2} \right) \quad (25.4)$$

When the average is calculated continuously, and r , q , and σ are constant (as in DerivaGem):

$$M_1 = \frac{e^{(r-q)T} - 1}{(r-q)T} S_0$$

and

$$M_2 = \frac{2e^{[2(r-q)+\sigma^2]T} S_0^2}{(r-q+\sigma^2)(2r-2q+\sigma^2)T^2} + \frac{2S_0^2}{(r-q)T^2} \left(\frac{1}{2(r-q)+\sigma^2} - \frac{e^{(r-q)T}}{r-q+\sigma^2} \right)$$

More generally, when the average is calculated from observations at times T_i ($1 \leq i \leq m$),

$$M_1 = \frac{1}{m} \sum_{i=1}^m F_i \quad \text{and} \quad M_2 = \frac{1}{m^2} \left(\sum_{i=1}^m F_i^2 e^{\sigma_i^2 T_i} + 2 \sum_{j=1}^m \sum_{i=1}^{j-1} F_i F_j e^{\sigma_i^2 T_i} \right)$$

where F_i and σ_i are the forward price and implied volatility for maturity T_i . See Technical Note 27 on www.rotman.utoronto.ca/~hull/TechnicalNotes for a proof of this.

Example 25.3

Consider a newly issued average price call option on a non-dividend-paying stock where the stock price is 50, the strike price is 50, the stock price volatility is 40% per annum, the risk-free rate is 10% per annum, and the time to maturity is 1 year. In this case, $S_0 = 50$, $K = 50$, $r = 0.1$, $q = 0$, $\sigma = 0.4$, and $T = 1$. If the average is calculated continuously, $M_1 = 52.59$ and $M_2 = 2,922.76$. From equations (25.3) and (25.4), $F_0 = 52.59$ and $\sigma = 23.54\%$. Equation (16.9), with $K = 50$, $T = 1$, and $r = 0.1$, gives the value of the option as 5.62. When 12, 52, and 250 observations are used for the average, the price is 6.00, 5.70, and 5.63, respectively.

¹⁰ When the asset price follows geometric Brownian motion, the geometric average of the price is exactly lognormal and the arithmetic average is approximately lognormal.

¹¹ See S. M. Turnbull and L. M. Wakeman, "A Quick Algorithm for Pricing European Average Options," *Journal of Financial and Quantitative Analysis*, 26 (September 1991): 377-89.

We can modify the analysis to accommodate the situation where the option is not newly issued and some prices used to determine the average have already been observed. Suppose that the averaging period is composed of a period of length t_1 over which prices have already been observed and a future period of length t_2 (the remaining life of the option). Suppose that the average asset price during the first time period is \bar{S} . The payoff from an average price call is

$$\max\left(\frac{\bar{S}t_1 + S_{\text{ave}}t_2}{t_1 + t_2} - K, 0\right)$$

where S_{ave} is the average asset price during the remaining part of the averaging period. This is the same as

$$\frac{t_2}{t_1 + t_2} \max(S_{\text{ave}} - K^*, 0)$$

where

$$K^* = \frac{t_1 + t_2}{t_2} K - \frac{t_1}{t_2} \bar{S}$$

When $K^* > 0$, the option can be valued in the same way as a newly issued Asian option provided that we change the strike price from K to K^* and multiply the result by $t_2/(t_1 + t_2)$. When $K^* < 0$ the option is certain to be exercised and can be valued as a forward contract. The value is

$$\frac{t_2}{t_1 + t_2} [M_1 e^{-rt_2} - K^* e^{-rt_2}]$$

Another type of Asian option is an average strike option. An *average strike call* pays off $\max(0, S_T - S_{\text{ave}})$ and an *average strike put* pays off $\max(0, S_{\text{ave}} - S_T)$. Average strike options can guarantee that the average price paid for an asset in frequent trading over a period of time is not greater than the final price. Alternatively, it can guarantee that the average price received for an asset in frequent trading over a period of time is not less than the final price. It can be valued as an option to exchange one asset for another when S_{ave} is assumed to be lognormal.

25.13 OPTIONS TO EXCHANGE ONE ASSET FOR ANOTHER

Options to exchange one asset for another (sometimes referred to as *exchange options*) arise in various contexts. An option to buy yen with Australian dollars is, from the point of view of a US investor, an option to exchange one foreign currency asset for another foreign currency asset. A stock tender offer is an option to exchange shares in one stock for shares in another stock.

Consider a European option to give up an asset worth U_T at time T and receive in return an asset worth V_T . The payoff from the option is

$$\max(V_T - U_T, 0)$$

A formula for valuing this option was first produced by Margrabe.¹² Suppose that the

¹² See W. Margrabe, "The Value of an Option to Exchange One Asset for Another," *Journal of Finance*, 33 (March 1978): 177–86.

asset prices, U and V , both follow geometric Brownian motion with volatilities σ_U and σ_V . Suppose further that the instantaneous correlation between U and V is ρ , and the yields provided by U and V are q_U and q_V , respectively. The value of the option at time zero is

$$V_0 e^{-q_V T} N(d_1) - U_0 e^{-q_U T} N(d_2) \quad (25.5)$$

where

$$d_1 = \frac{\ln(V_0/U_0) + (q_U - q_V + \hat{\sigma}^2/2)T}{\hat{\sigma}\sqrt{T}}, \quad d_2 = d_1 - \hat{\sigma}\sqrt{T}$$

and

$$\hat{\sigma} = \sqrt{\sigma_U^2 + \sigma_V^2 - 2\rho\sigma_U\sigma_V}$$

and U_0 and V_0 are the values of U and V at times zero.

This result will be proved in Chapter 27. It is interesting to note that equation (25.5) is independent of the risk-free rate r . This is because, as r increases, the growth rate of both asset prices in a risk-neutral world increases, but this is exactly offset by an increase in the discount rate. The variable $\hat{\sigma}$ is the volatility of V/U . Comparisons with equation (16.4) show that the option price is the same as the price of U_0 European call options on an asset worth V/U when the strike price is 1.0, the risk-free interest rate is q_U , and the dividend yield on the asset is q_V . Mark Rubinstein shows that the American version of this option can be characterized similarly for valuation purposes.¹³ It can be regarded as U_0 American options to buy an asset worth V/U for 1.0 when the risk-free interest rate is q_U and the dividend yield on the asset is q_V . The option can therefore be valued as described in Chapter 20 using a binomial tree.

An option to obtain the better or worse of two assets can be regarded as a position in one of the assets combined with an option to exchange it for the other asset:

$$\min(U_T, V_T) = V_T - \max(V_T - U_T, 0)$$

$$\max(U_T, V_T) = U_T + \max(V_T - U_T, 0)$$

25.14 OPTIONS INVOLVING SEVERAL ASSETS

Options involving two or more risky assets are sometimes referred to as *rainbow options*. One example is the bond futures contract traded on the CBOT described in Chapter 6. The party with the short position is allowed to choose between a large number of different bonds when making delivery.

Probably the most popular option involving several assets is a European *basket option*. This is an option where the payoff is dependent on the value of a portfolio (or basket) of assets. The assets are usually either individual stocks or stock indices or currencies. A European basket option can be valued with Monte Carlo simulation, by assuming that the assets follow correlated geometric Brownian motion processes. A much faster approach is to calculate the first two moments of the basket at the maturity of the option in a risk-neutral world, and then assume that value of the basket is

¹³ See M. Rubinstein, "One for Another," *Risk*, July/August 1991: 30–32

lognormally distributed at that time. The option can then be valued using Black's model with the parameters shown in equations (25.3) and (25.4). In this case,

$$M_1 = \sum_{i=1}^n F_i \quad \text{and} \quad M_2 = \sum_{i=1}^n \sum_{j=1}^n F_i F_j e^{\rho_{ij} \sigma_i \sigma_j T}$$

where n is the number of assets, T is the option maturity, F_i and σ_i are the forward price and volatility of the i th asset, and ρ_{ij} is the correlation between the i th and j th asset. See Technical Note 28 at www.rotman.utoronto.ca/~hull/TechnicalNotes.

25.15 VOLATILITY AND VARIANCE SWAPS

A volatility swap is an agreement to exchange the realized volatility of an asset between time 0 and time T for a prespecified fixed volatility. The realized volatility is usually calculated as described in Section 14.4 but with the assumption that the mean daily return is zero. Suppose that there are n daily observations on the asset price during the period between time 0 and time T . The realized volatility is

$$\bar{\sigma} = \sqrt{\frac{252}{n-2} \sum_{i=1}^{n-1} \left[\ln \left(\frac{S_{i+1}}{S_i} \right) \right]^2}$$

where S_i is the i th observation on the asset price. (Sometimes $n-1$ might replace $n-2$ in this formula.)

The payoff from the volatility swap at time T to the payer of the fixed volatility is $L_{\text{vol}}(\bar{\sigma} - \sigma_K)$, where L_{vol} is the notional principal and σ_K is the fixed volatility. Whereas an option provides a complex exposure to the asset price and volatility, a volatility swap is simpler in that it has exposure only to volatility.

A variance swap is an agreement to exchange the realized variance rate \bar{V} between time 0 and time T for a prespecified variance rate. The variance rate is the square of the volatility ($\bar{V} = \bar{\sigma}^2$). Variance swaps are easier to value than volatility swaps. This is because the variance rate between time 0 and time T can be replicated using a portfolio of put and call options. The payoff from a variance swap at time T to the payer of the fixed variance rate is $L_{\text{var}}(\bar{V} - V_K)$, where L_{var} is the notional principal and V_K is the fixed variance rate. Often the notional principal for a variance swap is expressed in terms of the corresponding notional principal for a volatility swap using $L_{\text{var}} = L_{\text{vol}}/(2\sigma_K)$.

Valuation of Variance Swap

Technical Note 22 at www.rotman.utoronto.ca/~hull/TechnicalNotes shows that, for any value S^* of the asset price, the expected average variance between times 0 and T is

$$\hat{E}(\bar{V}) = \frac{2}{T} \ln \frac{F_0}{S^*} - \frac{2}{T} \left[\frac{F_0}{S^*} - 1 \right] + \frac{2}{T} \left[\int_{K=0}^{S^*} \frac{1}{K^2} e^{rT} p(K) dK + \int_{K=S^*}^{\infty} \frac{1}{K^2} e^{rT} c(K) dK \right] \quad (25.6)$$

where F_0 is the forward price of the asset for a contract maturing at time T , $c(K)$ is

the price of a European call option with strike price K and time to maturity T , and $p(K)$ is the price of a European put option with strike price K and time to maturity T .

This provides a way of valuing a variance swap.¹⁴ The value of an agreement to receive the realized variance between time 0 and time T and pay a variance rate of V_K , with both being applied to a principal of L_{var} , is

$$L_{\text{var}}[\hat{E}(\bar{V}) - V_K]e^{-rT} \quad (25.7)$$

Suppose that the prices of European options with strike prices K_i ($1 \leq i \leq n$) are known, where $K_1 < K_2 < \dots < K_n$. A standard approach for implementing equation (25.6) is to set S^* equal to the first strike price below F_0 and then approximate the integrals as

$$\int_{K=0}^{S^*} \frac{1}{K^2} e^{rT} p(K) dK + \int_{K=S^*}^{\infty} \frac{1}{K^2} e^{rT} c(K) dK = \sum_{i=1}^n \frac{\Delta K_i}{K_i^2} e^{rT} Q(K_i) \quad (25.8)$$

where $\Delta K_i = 0.5(K_{i+1} - K_{i-1})$ for $2 \leq i \leq n-1$, $\Delta K_1 = K_2 - K_1$, $\Delta K_n = K_n - K_{n-1}$. The function $Q(K_i)$ is the price of a European put option with strike price K_i if $K_i < S^*$ and the price of a European call option with strike price K_i if $K_i > S^*$. When $K_i = S^*$, the function $Q(K_i)$ is equal to the average of the prices of a European call and a European put with strike price K_i .

Example 25.4

Consider a 3-month contract to receive the realized variance rate of an index over the 3 months and pay a variance rate of 0.045 on a principal of \$100 million. The risk-free rate is 4% and the dividend yield on the index is 1%. The current level of the index is 1020. Suppose that, for strike prices of 800, 850, 900, 950, 1,000, 1,050, 1,100, 1,150, 1,200, the 3-month implied volatilities of the index are 29%, 28%, 27%, 26%, 25%, 24%, 23%, 22%, 21%, respectively. In this case, $n = 9$, $K_1 = 800$, $K_2 = 850, \dots, K_9 = 1,200$, $F_0 = 1,020e^{(0.04-0.01) \times 0.25} = 1,027.68$, and $S^* = 1,000$. DerivaGem shows that $Q(K_1) = 2.22$, $Q(K_2) = 5.22$, $Q(K_3) = 11.05$, $Q(K_4) = 21.27$, $Q(K_5) = 51.21$, $Q(K_6) = 38.94$, $Q(K_7) = 20.69$, $Q(K_8) = 9.44$, $Q(K_9) = 3.57$. Also, $\Delta K_i = 50$ for all i . Hence,

$$\sum_i^n \frac{\Delta K_i}{K_i^2} e^{rT} Q(K_i) = 0.008139$$

From equations (25.6) and (25.8), it follows that

$$\hat{E}(\bar{V}) = \frac{2}{0.25} \ln\left(\frac{1027.68}{1,000}\right) - \frac{2}{0.25} \left(\frac{1027.68}{1,000} - 1\right) + \frac{2}{0.25} \times 0.008139 = 0.0621$$

From equation (25.7), the value of the variance swap (in millions of dollars) is $100 \times (0.0621 - 0.045)e^{-0.04 \times 0.25} = 1.69$.

¹⁴ See also K. Demeterfi, E. Derman, M. Kamal, and J. Zou, "A Guide to Volatility and Variance Swaps," *The Journal of Derivatives*, 6, 4 (Summer 1999), 9–32. For options on variance and volatility, see P. Carr and R. Lee, "Realized Volatility and Variance: Options via Swaps," *Risk*, May 2007, 76–83.

Valuation of a Volatility Swap

To value a volatility swap, we require $\hat{E}(\bar{\sigma})$, where $\bar{\sigma}$ is the average value of volatility between time 0 and time T . We can write

$$\bar{\sigma} = \sqrt{\hat{E}(\bar{V})} \sqrt{1 + \frac{\bar{V} - \hat{E}(\bar{V})}{\hat{E}(\bar{V})}}$$

Expanding the second term on the right-hand side in a series gives

$$\bar{\sigma} = \sqrt{\hat{E}(\bar{V})} \left\{ 1 + \frac{\bar{V} - \hat{E}(\bar{V})}{2\hat{E}(\bar{V})} - \frac{1}{8} \left[\frac{\bar{V} - \hat{E}(\bar{V})}{\hat{E}(\bar{V})} \right]^2 \right\}$$

Taking expectations,

$$\hat{E}(\bar{\sigma}) = \sqrt{\hat{E}(\bar{V})} \left\{ 1 - \frac{1}{8} \left[\frac{\text{var}(\bar{V})}{\hat{E}(\bar{V})^2} \right] \right\} \quad (25.9)$$

where $\text{var}(\bar{V})$ is the variance of \bar{V} . The valuation of a volatility swap therefore requires an estimate of the variance of the average variance rate during the life of the contract. The value of an agreement to receive the realized volatility between time 0 and time T and pay a volatility of σ_K , with both being applied to a principal of L_{vol} , is

$$L_{\text{vol}}[\hat{E}(\bar{\sigma}) - \sigma_K]e^{-rT}$$

Example 25.5

For the situation in Example 25.4, consider a volatility swap where the realized volatility is received and a volatility of 23% is paid on a principal of \$100 million. In this case $\hat{E}(\bar{V}) = 0.0621$. Suppose that the standard deviation of the average variance over 3 months has been estimated as 0.01. This means that $\text{var}(\bar{V}) = 0.0001$. Equation (25.9) gives

$$\hat{E}(\bar{\sigma}) = \sqrt{0.0621} \left(1 - \frac{1}{8} \times \frac{0.0001}{0.0621^2} \right) = 0.2484$$

The value of the swap in (millions of dollars) is

$$100 \times (0.2484 - 0.23)e^{-0.04 \times 0.25} = 1.82$$

The VIX Index

In equation (25.6), the \ln function can be approximated by the first two terms in a series expansion:

$$\ln\left(\frac{F_0}{S^*}\right) = \left(\frac{F_0}{S^*} - 1\right) - \frac{1}{2}\left(\frac{F_0}{S^*} - 1\right)^2$$

This means that the risk-neutral expected cumulative variance is calculated as

$$\hat{E}(\bar{V})T = -\left(\frac{F_0}{S^*} - 1\right)^2 + 2 \sum_{i=1}^n \frac{\Delta K_i}{K_i^2} e^{rT} Q(K_i) \quad (25.10)$$

Since 2004 the VIX volatility index (see Section 14.11) has been based on equation (25.10). The procedure used on any given day is to calculate $\hat{E}(\bar{V})T$ for options that

trade in the market and have maturities immediately above and below 30 days. The 30-day risk-neutral expected cumulative variance is calculated from these two numbers using interpolation. This is then multiplied by 365/30 and the index is set equal to the square root of the result. More details on the calculation can be found on:

www.cboe.com/micro/vix/vixwhite.pdf

25.16 STATIC OPTIONS REPLICATION

If the procedures described in Chapter 18 are used for hedging exotic options, some are easy to handle, but others are very difficult because of discontinuities (see Business Snapshot 25.1). For the difficult cases, a technique known as static options replication is sometimes useful.¹⁵ This involves searching for a portfolio of actively traded options that approximately replicates the exotic option. Shorting this position provides the hedge.¹⁶

The basic principle underlying static options replication is as follows. If two portfolios are worth the same on a certain boundary, they are also worth the same at all interior points of the boundary. Consider as an example a 9-month up-and-out call option on a non-dividend-paying stock where the stock price is 50, the strike price is 50, the barrier is 60, the risk-free interest rate is 10% per annum, and the volatility is 30% per annum. Suppose that $f(S, t)$ is the value of the option at time t for a stock price of S . Any boundary in (S, t) space can be used for the purposes of producing the replicating portfolio. A convenient one to choose is shown in Figure 25.1. It is defined by $S = 60$ and $t = 0.75$. The values of the up-and-out option on the boundary are given by

$$f(S, 0.75) = \max(S - 50, 0) \quad \text{when } S < 60$$

$$f(60, t) = 0 \quad \text{when } 0 \leq t \leq 0.75$$

There are many ways that these boundary values can be approximately matched using regular options. The natural option to match the first boundary is a 9-month European call with a strike price of 50. The first component of the replicating portfolio is therefore one unit of this option. (We refer to this option as option A.)

One way of matching the $f(60, t)$ boundary is to proceed as follows:

1. Divide the life of the option into N steps of length Δt
2. Choose a European call option with a strike price of 60 and maturity at time $N\Delta t$ (= 9 months) to match the boundary at the $\{60, (N - 1)\Delta t\}$ point
3. Choose a European call option with a strike price of 60 and maturity at time $(N - 1)\Delta t$ to match the boundary at the $\{60, (N - 2)\Delta t\}$ point

and so on. Note that the options are chosen in sequence so that they have zero value on the parts of the boundary matched by earlier options.¹⁷ The option with a strike price

¹⁵ See E. Derman, D. Ergener, and I. Kani, "Static Options Replication," *Journal of Derivatives* 2, 4 (Summer 1995): 78–95.

¹⁶ Technical Note 22 at www.rotman.utoronto.ca/~hull/TechnicalNotes provides an example of static replication. It shows that the variance rate of an asset can be replicated by a position in the asset and out-of-the-money options on the asset. This result, which leads to equation (25.6), can be used to hedge variance swaps.

¹⁷ This is not a requirement. If K points on the boundary are to be matched, we can choose K options and solve a set of K linear equations to determine required positions in the options.

Business Snapshot 25.1 Is Delta Hedging Easier or More Difficult for Exotics?

As described in Chapter 18 we can approach the hedging of exotic options by creating a delta neutral position and rebalancing frequently to maintain delta neutrality. When we do this we find some exotic options are easier to hedge than plain vanilla options and some are more difficult.

An example of an exotic option that is relatively easy to hedge is an average price option where the averaging period is the whole life of the option. As time passes, we observe more of the asset prices that will be used in calculating the final average. This means that our uncertainty about the payoff decreases with the passage of time. As a result, the option becomes progressively easier to hedge. In the final few days, the delta of the option always approaches zero because price movements during this time have very little impact on the payoff.

By contrast barrier options are relatively difficult to hedge. Consider a down-and-out call option on a currency when the exchange rate is 0.0005 above the barrier. If the barrier is hit, the option is worth nothing. If the barrier is not hit, the option may prove to be quite valuable. The delta of the option is discontinuous at the barrier making conventional hedging very difficult.

of 60 that matures in 9 months has zero value on the vertical boundary that is matched by option A. The option maturing at time $i \Delta t$ has zero value at the point $\{60, i \Delta t\}$ that is matched by the option maturing at time $(i + 1)\Delta t$ for $1 \leq i \leq N - 1$.

Suppose that $\Delta = 0.25$. In addition to option A, the replicating portfolio consists of positions in European options with strike price 60 that mature in 9, 6, and 3 months. We will refer to these as options B, C, and D, respectively. Given our assumptions

Figure 25.1 Boundary points used for static options replication example.

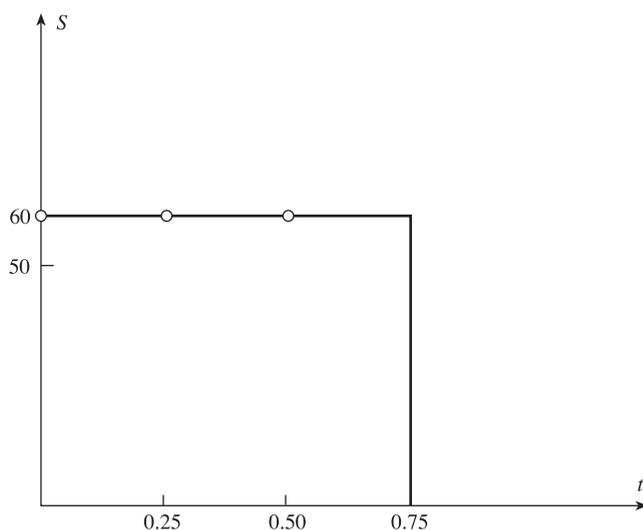


Table 25.1 The portfolio of European call options used to replicate an up-and-out option.

<i>Option</i>	<i>Strike price</i>	<i>Maturity (years)</i>	<i>Position</i>	<i>Initial value</i>
A	50	0.75	1.00	+6.99
B	60	0.75	-2.66	-8.21
C	60	0.50	0.97	+1.78
D	60	0.25	0.28	+0.17

about volatility and interest rates, option B is worth 4.33 at the {60, 0.5} point. Option A is worth 11.54 at this point. The position in option B necessary to match the boundary at the {60, 0.5} point is therefore $-11.54/4.33 = -2.66$. Option C is worth 4.33 at the {60, 0.25} point. The position taken in options A and B is worth -4.21 at this point. The position in option C necessary to match the boundary at the {60, 0.25} point is therefore $4.21/4.33 = 0.97$. Similar calculations show that the position in option D necessary to match the boundary at the {60, 0} point is 0.28.

The portfolio chosen is summarized in Table 25.1. (See also Sample Application F of the DerivaGem Applications.) It is worth 0.73 initially (i.e., at time zero when the stock price is 50). This compares with 0.31 given by the analytic formula for the up-and-out call earlier in this chapter. The replicating portfolio is not exactly the same as the up-and-out option because it matches the latter at only three points on the second boundary. If we use the same procedure, but match at 18 points on the second boundary (using options that mature every half month), the value of the replicating portfolio reduces to 0.38. If 100 points are matched, the value reduces further to 0.32.

To hedge a derivative, the portfolio that replicates its boundary conditions must be shorted. The portfolio must be unwound when any part of the boundary is reached.

Static options replication has the advantage over delta hedging that it does not require frequent rebalancing. It can be used for a wide range of derivatives. The user has a great deal of flexibility in choosing the boundary that is to be matched and the options that are to be used.

SUMMARY

Exotic options are options with rules governing the payoff that are more complicated than standard options. We have discussed 14 different types of exotic options: packages, nonstandard American options, gap options, forward start options, cliquet options, compound options, chooser options, barrier options, binary options, lookback options, shout options, Asian options, options to exchange one asset for another, and options involving several assets. We have discussed how these can be valued using the same assumptions as those used to derive the Black–Scholes–Merton model in Chapter 14. Some can be valued analytically, but using much more complicated formulas than those for regular European calls and puts, some can be handled using analytic approximations, and some can be valued using extensions of the numerical procedures in Chapter 20. We will present more numerical procedures for valuing exotic options in Chapter 26.

Some exotic options are easier to hedge than the corresponding regular options; others are more difficult. In general, Asian options are easier to hedge because the payoff becomes progressively more certain as we approach maturity. Barrier options can be more difficult to hedge because delta is discontinuous at the barrier. One approach to hedging an exotic option, known as static options replication, is to find a portfolio of regular options whose value matches the value of the exotic option on some boundary. The exotic option is hedged by shorting this portfolio.

FURTHER READING

- Carr, P., and R. Lee, "Realized Volatility and Variance: Options via Swaps," *Risk*, May 2007, 76–83.
- Clelow, L., and C. Strickland, *Exotic Options: The State of the Art*. London: Thomson Business Press, 1997.
- Demeterfi, K., E. Derman, M. Kamal, and J. Zou, "More than You Ever Wanted to Know about Volatility Swaps," *Journal of Derivatives*, 6, 4 (Summer, 1999), 9–32.
- Derman, E., D. Ergener, and I. Kani, "Static Options Replication," *Journal of Derivatives*, 2, 4 (Summer 1995): 78–95.
- Geske, R., "The Valuation of Compound Options," *Journal of Financial Economics*, 7 (1979): 63–81.
- Goldman, B., H. Sosin, and M. A. Gatto, "Path Dependent Options: Buy at the Low, Sell at the High," *Journal of Finance*, 34 (December 1979): 1111–27.
- Margrabe, W., "The Value of an Option to Exchange One Asset for Another," *Journal of Finance*, 33 (March 1978): 177–86.
- Milevsky, M. A., and S. E. Posner, "Asian Options: The Sum of Lognormals and the Reciprocal Gamma Distribution," *Journal of Financial and Quantitative Analysis*, 33, 3 (September 1998), 409–22.
- Ritchken, P. "On Pricing Barrier Options," *Journal of Derivatives*, 3, 2 (Winter 1995): 19–28.
- Ritchken P., L. Sankarasubramanian, and A. M. Vihj, "The Valuation of Path Dependent Contracts on the Average," *Management Science*, 39 (1993): 1202–13.
- Rubinstein, M., and E. Reiner, "Breaking Down the Barriers," *Risk*, September (1991): 28–35.
- Rubinstein, M., "Double Trouble," *Risk*, December/January (1991/1992): 53–56.
- Rubinstein, M., "One for Another," *Risk*, July/August (1991): 30–32.
- Rubinstein, M., "Options for the Undecided," *Risk*, April (1991): 70–73.
- Rubinstein, M., "Pay Now, Choose Later," *Risk*, February (1991): 44–47.
- Rubinstein, M., "Somewhere Over the Rainbow," *Risk*, November (1991): 63–66.
- Rubinstein, M., "Two in One," *Risk* May (1991): 49.
- Rubinstein, M., and E. Reiner, "Unscrambling the Binary Code," *Risk*, October 1991: 75–83.
- Stulz, R. M., "Options on the Minimum or Maximum of Two Assets," *Journal of Financial Economics*, 10 (1982): 161–85.
- Turnbull, S. M., and L. M. Wakeman, "A Quick Algorithm for Pricing European Average Options," *Journal of Financial and Quantitative Analysis*, 26 (September 1991): 377–89.

Practice Questions (Answers in Solutions Manual)

- 25.1. Explain the difference between a forward start option and a chooser option.
- 25.2. Describe the payoff from a portfolio consisting of a floating lookback call and a floating lookback put with the same maturity.

- 25.3. Consider a chooser option where the holder has the right to choose between a European call and a European put at any time during a 2-year period. The maturity dates and strike prices for the calls and puts are the same regardless of when the choice is made. Is it ever optimal to make the choice before the end of the 2-year period? Explain your answer.
- 25.4. Suppose that c_1 and p_1 are the prices of a European average price call and a European average price put with strike price K and maturity T , c_2 and p_2 are the prices of a European average strike call and European average strike put with maturity T , and c_3 and p_3 are the prices of a regular European call and a regular European put with strike price K and maturity T . Show that $c_1 + c_2 - c_3 = p_1 + p_2 - p_3$.
- 25.5. The text derives a decomposition of a particular type of chooser option into a call maturing at time T_2 and a put maturing at time T_1 . Derive an alternative decomposition into a call maturing at time T_1 and a put maturing at time T_2 .
- 25.6. Section 25.8 gives two formulas for a down-and-out call. The first applies to the situation where the barrier, H , is less than or equal to the strike price, K . The second applies to the situation where $H \geq K$. Show that the two formulas are the same when $H = K$.
- 25.7. Explain why a down-and-out put is worth zero when the barrier is greater than the strike price.
- 25.8. Suppose that the strike price of an American call option on a non-dividend-paying stock grows at rate g . Show that if g is less than the risk-free rate, r , it is never optimal to exercise the call early.
- 25.9. How can the value of a forward start put option on a non-dividend-paying stock be calculated if it is agreed that the strike price will be 10% greater than the stock price at the time the option starts?
- 25.10. If a stock price follows geometric Brownian motion, what process does $A(t)$ follow where $A(t)$ is the arithmetic average stock price between time zero and time t ?
- 25.11. Explain why delta hedging is easier for Asian options than for regular options.
- 25.12. Calculate the price of a 1-year European option to give up 100 ounces of silver in exchange for 1 ounce of gold. The current prices of gold and silver are \$380 and \$4, respectively; the risk-free interest rate is 10% per annum; the volatility of each commodity price is 20%; and the correlation between the two prices is 0.7. Ignore storage costs.
- 25.13. Is a European down-and-out option on an asset worth the same as a European down-and-out option on the asset's futures price for a futures contract maturing at the same time as the option?
- 25.14. Answer the following questions about compound options:
- What put-call parity relationship exists between the price of a European call on a call and a European put on a call? Show that the formulas given in the text satisfy the relationship.
 - What put-call parity relationship exists between the price of a European call on a put and a European put on a put? Show that the formulas given in the text satisfy the relationship.
- 25.15. Does a floating lookback call become more valuable or less valuable as we increase the frequency with which we observe the asset price in calculating the minimum?

- 25.16. Does a down-and-out call become more valuable or less valuable as we increase the frequency with which we observe the asset price in determining whether the barrier has been crossed? What is the answer to the same question for a down-and-in call?
- 25.17. Explain why a regular European call option is the sum of a down-and-out European call and a down-and-in European call. Is the same true for American call options?
- 25.18. What is the value of a derivative that pays off \$100 in 6 months if the S&P 500 index is greater than 1,000 and zero otherwise? Assume that the current level of the index is 960, the risk-free rate is 8% per annum, the dividend yield on the index is 3% per annum, and the volatility of the index is 20%.
- 25.19. In a 3-month down-and-out call option on silver futures the strike price is \$20 per ounce and the barrier is \$18. The current futures price is \$19, the risk-free interest rate is 5%, and the volatility of silver futures is 40% per annum. Explain how the option works and calculate its value. What is the value of a regular call option on silver futures with the same terms? What is the value of a down-and-in call option on silver futures with the same terms?
- 25.20. A new European-style floating lookback call option on a stock index has a maturity of 9 months. The current level of the index is 400, the risk-free rate is 6% per annum, the dividend yield on the index is 4% per annum, and the volatility of the index is 20%. Use DerivaGem to value the option.
- 25.21. Estimate the value of a new 6-month European-style average price call option on a non-dividend-paying stock. The initial stock price is \$30, the strike price is \$30, the risk-free interest rate is 5%, and the stock price volatility is 30%.
- 25.22. Use DerivaGem to calculate the value of:
- A regular European call option on a non-dividend-paying stock where the stock price is \$50, the strike price is \$50, the risk-free rate is 5% per annum, the volatility is 30%, and the time to maturity is one year
 - A down-and-out European call which is as in (a) with the barrier at \$45
 - A down-and-in European call which is as in (a) with the barrier at \$45.
- Show that the option in (a) is worth the sum of the values of the options in (b) and (c).
- 25.23. Explain adjustments that have to be made when $r = q$ for (a) the valuation formulas for floating lookback call options in Section 25.10 and (b) the formulas for M_1 and M_2 in Section 25.12.
- 25.24. Value the variance swap in Example 25.4 of Section 25.15 assuming that the implied volatilities for options with strike prices 800, 850, 900, 950, 1,000, 1,050, 1,100, 1,150, 1,200 are 20%, 20.5%, 21%, 21.5%, 22%, 22.5%, 23%, 23.5%, 24%, respectively.

Further Questions

- 25.25. What is the value in dollars of a derivative that pays off £10,000 in 1 year provided that the dollar/sterling exchange rate is greater than 1.5000 at that time? The current exchange rate is 1.4800. The dollar and sterling interest rates are 4% and 8% per annum, respectively. The volatility of the exchange rate is 12% per annum.
- 25.26. Consider an up-and-out barrier call option on a non-dividend-paying stock when the stock price is 50, the strike price is 50, the volatility is 30%, the risk-free rate is 5%, the

- time to maturity is 1 year, and the barrier at \$80. Use the software to value the option and graph the relationship between (a) the option price and the stock price, (b) the delta and the stock price, (c) the option price and the time to maturity, and (d) the option price and the volatility. Provide an intuitive explanation for the results you get. Show that the delta, gamma, theta, and vega for an up-and-out barrier call option can be either positive or negative.
- 25.27. Sample Application F in the DerivaGem Application Builder Software considers the static options replication example in Section 25.15. It shows the way a hedge can be constructed using four options (as in Section 25.15) and two ways a hedge can be constructed using 16 options.
- Explain the difference between the two ways a hedge can be constructed using 16 options. Explain intuitively why the second method works better.
 - Improve on the four-option hedge by changing T_{mat} for the third and fourth options.
 - Check how well the 16-option portfolios match the delta, gamma, and vega of the barrier option.
- 25.28. Consider a down-and-out call option on a foreign currency. The initial exchange rate is 0.90, the time to maturity is 2 years, the strike price is 1.00, the barrier is 0.80, the domestic risk-free interest rate is 5%, the foreign risk-free interest rate is 6%, and the volatility is 25% per annum. Use DerivaGem to develop a static option replication strategy involving five options.
- 25.29. Suppose that a stock index is currently 900. The dividend yield is 2%, the risk-free rate is 5%, and the volatility is 40%. Use the results in the appendix to calculate the value of a 1-year average price call where the strike price is 900 and the index level is observed at the end of each quarter for the purposes of the averaging. Compare this with the price calculated by DerivaGem for a 1-year average price option where the price is observed continuously. Provide an intuitive explanation for any differences between the prices.
- 25.30. Use the DerivaGem Application Builder software to compare the effectiveness of daily delta hedging for (a) the option considered in Tables 18.2 and 18.3 and (b) an average price call with the same parameters. Use Sample Application C. For the average price option you will find it necessary to change the calculation of the option price in cell C16, the payoffs in cells H15 and H16, and the deltas (cells G46 to G186 and N46 to N186). Carry out 20 Monte Carlo simulation runs for each option by repeatedly pressing F9. On each run record the cost of writing and hedging the option, the volume of trading over the whole 20 weeks and the volume of trading between weeks 11 and 20. Comment on the results.
- 25.31. In the DerivaGem Application Builder Software modify Sample Application D to test the effectiveness of delta and gamma hedging for a call on call compound option on a 100,000 units of a foreign currency where the exchange rate is 0.67, the domestic risk-free rate is 5%, the foreign risk-free rate is 6%, the volatility is 12%. The time to maturity of the first option is 20 weeks, and the strike price of the first option is 0.015. The second option matures 40 weeks from today and has a strike price of 0.68. Explain how you modified the cells. Comment on hedge effectiveness.
- 25.32. Outperformance certificates (also called “sprint certificates,” “accelerator certificates,” or “speeders”) are offered to investors by many European banks as a way of investing in a company’s stock. The initial investment equals the stock price, S_0 . If the stock price goes up between time 0 and time T , the investor gains k times the increase at time T ,

where k is a constant greater than 1.0. However, the stock price used to calculate the gain at time T is capped at some maximum level M . If the stock price goes down, the investor's loss is equal to the decrease. The investor does not receive dividends.

- (a) Show that an outperformance certificate is a package.
 - (b) Calculate using DerivaGem the value of a one-year outperformance certificate when the stock price is 50 euros, $k = 1.5$, $M = 70$ euros, the risk-free rate is 5%, and the stock price volatility is 25%. Dividends equal to 0.5 euros are expected in 2 months, 5 months, 8 months, and 11 months.
- 25.33. Carry out the analysis in Example 25.4 of Section 25.15 to value the variance swap on the assumption that the life of the swap is 1 month rather than 3 months.
- 25.34. What is the relationship between a regular call option, a binary call option, and a gap call option?
- 25.35. Produce a formula for valuing a cliquet option where an amount Q is invested to produce a payoff at the end of n periods. The return earned each period is the greater of the return on an index (excluding dividends) and zero.



26

CHAPTER

More on Models and Numerical Procedures

Up to now the models we have used to value options have been based on the geometric Brownian motion model of asset price behavior that underlies the Black–Scholes–Merton formulas and the numerical procedures we have used have been relatively straightforward. In this chapter we introduce a number of new models and explain how the numerical procedures can be adapted to cope with particular situations.

Chapter 19 explained how traders overcome the weaknesses in the geometric Brownian motion model by using volatility surfaces. A volatility surface determines an appropriate volatility to substitute into Black–Scholes–Merton when pricing plain vanilla options. Unfortunately it says little about the volatility that should be used for exotic options when the pricing formulas of Chapter 25 are used. Suppose the volatility surface shows that the correct volatility to use when pricing a one-year plain vanilla option with a strike price of \$40 is 27%. This is liable to be totally inappropriate for pricing a barrier option (or some other exotic option) that has a strike price of \$40 and a life of one year.

The first part of this chapter discusses a number of alternatives to geometric Brownian motion that are designed to deal with the problem of pricing exotic options consistently with plain vanilla options. These alternative asset price processes fit the market prices of plain vanilla options better than geometric Brownian motion. As a result, we can have more confidence in using them to value exotic options.

The second part of the chapter extends the discussion of numerical procedures. It explains how convertible bonds and some types of path-dependent derivatives can be valued using trees. It discusses the special problems associated with valuing barrier options numerically and how these problems can be handled. Finally, it outlines alternative ways of constructing trees for two correlated variables and shows how Monte Carlo simulation can be used to value derivatives when there are early exercise opportunities.

As in earlier chapters, results are presented for derivatives dependent on an asset providing a yield at rate q . For an option on a stock index q should be set equal to the dividend yield on the index, for an option on a currency it should be set equal to the foreign risk-free rate, and for an option on a futures contract it should be set equal to the domestic risk-free rate.

26.1 ALTERNATIVES TO BLACK–SCHOLES–MERTON

The Black–Scholes–Merton model assumes that an asset’s price changes continuously in a way that produces a lognormal distribution for the price at any future time. There are many alternative processes that can be assumed. One possibility is to retain the property that the asset price changes continuously, but assume a process other than geometric Brownian motion. Another alternative is to overlay continuous asset price changes with jumps. Yet another alternative is to assume a process where all the asset price changes that take place are jumps. We will consider examples of all three types of processes in this section. A model where stock prices change continuously is known as a *diffusion model*. A model where continuous changes are overlaid with jumps is known as a *mixed jump–diffusion model*. A model where all stock price changes are jumps is known as a *pure jump model*. These types of processes are known collectively as *Levy processes*.¹

The Constant Elasticity of Variance Model

One alternative to Black–Scholes–Merton is the *constant elasticity of variance* (CEV) model. This is a diffusion model where the risk-neutral process for a stock price S is

$$dS = (r - q)S dt + \sigma S^\alpha dz$$

where r is the risk-free rate, q is the dividend yield, dz is a Wiener process, σ is a volatility parameter, and α is a positive constant.²

When $\alpha = 1$, the CEV model is the geometric Brownian motion model we have been using up to now. When $\alpha < 1$, the volatility increases as the stock price decreases. This creates a probability distribution similar to that observed for equities with a heavy left tail and less heavy right tail (see Figure 19.4).³ When $\alpha > 1$, the volatility increases as the stock price increases. This creates a probability distribution with a heavy right tail and a less heavy left tail. This corresponds to a volatility smile where the implied volatility is an increasing function of the strike price. This type of volatility smile is sometimes observed for options on futures (see Problem 17.23).

The valuation formulas for European call and put options under the CEV model are

$$c = S_0 e^{-qT} [1 - \chi^2(a, b + 2, c)] - K e^{-rT} \chi^2(c, b, a)$$

$$p = K e^{-rT} [1 - \chi^2(c, b, a)] - S_0 e^{-qT} \chi^2(a, b + 2, c)$$

when $0 < \alpha < 1$, and

$$c = S_0 e^{-qT} [1 - \chi^2(c, -b, a)] - K e^{-rT} \chi^2(a, 2 - b, c)$$

$$p = K e^{-rT} [1 - \chi^2(a, 2 - b, c)] - S_0 e^{-qT} \chi^2(c, -b, a)$$

¹ Roughly speaking, a Levy process is a continuous-time stochastic process with stationary independent increments.

² See J. C. Cox and S. A. Ross, “The Valuation of Options for Alternative Stochastic Processes,” *Journal of Financial Economics*, 3 (March 1976): 145–66.

³ The reason is as follows. As the stock price decreases, the volatility increases making even lower stock price more likely; when the stock price increases, the volatility decreases making higher stock prices less likely.

when $\alpha > 1$, with

$$a = \frac{[Ke^{-(r-q)T}]^{2(1-\alpha)}}{(1-\alpha)^2 v}, \quad b = \frac{1}{1-\alpha}, \quad c = \frac{S^{2(1-\alpha)}}{(1-\alpha)^2 v}$$

where

$$v = \frac{\sigma^2}{2(r-q)(\alpha-1)} [e^{2(r-q)(\alpha-1)T} - 1]$$

and $\chi^2(z, k, v)$ is the cumulative probability that a variable with a noncentral χ^2 distribution with noncentrality parameter v and k degrees of freedom is less than z . A procedure for computing $\chi^2(z, k, v)$ is provided in Technical Note 12 on the author's website: www.rotman.utoronto.ca/~hull/TechnicalNotes.

The CEV model is particularly useful for valuing exotic equity options. The parameters of the model can be chosen to fit the prices of plain vanilla options as closely as possible by minimizing the sum of the squared differences between model prices and market prices.

Merton's Mixed Jump-Diffusion Model

Merton has suggested a model where jumps are combined with continuous changes.⁴ Define:

λ : Average number of jumps per year

k : Average jump size measured as a percentage of the asset price

The percentage jump size is assumed to be drawn from a probability distribution in the model.

The probability of a jump in time Δt is $\lambda \Delta t$. The average growth rate in the asset price from the jumps is therefore λk . The risk-neutral process for the asset price is

$$\frac{dS}{S} = (r - q - \lambda k) dt + \sigma dz + dp$$

where dz is a Wiener process, dp is the Poisson process generating the jumps, and σ is the volatility of the geometric Brownian motion. The processes dz and dp are assumed to be independent.

An important particular case of Merton's model is where the logarithm of the size of the percentage jump is normal. Assume that the standard deviation of the normal distribution is s . Merton shows that a European option price can then be written

$$\sum_{n=0}^{\infty} \frac{e^{-\lambda'T} (\lambda'T)^n}{n!} f_n$$

where $\lambda' = \lambda(1+k)$. The variable f_n is the Black-Scholes-Merton option price when the dividend yield is q , the variance rate is

$$\sigma^2 + \frac{ns^2}{T}$$

⁴ See R. C. Merton, "Option Pricing When Underlying Stock Returns Are Discontinuous," *Journal of Financial Economics*, 3 (March 1976): 125-44.

and the risk-free rate is

$$r - \lambda k + \frac{n\gamma}{T}$$

where $\gamma = \ln(1 + k)$.

This model gives rise to heavier left and heavier right tails than Black–Scholes–Merton. It can be used for pricing currency options. As in the case of the CEV model, the model parameters are chosen by minimizing the sum of the squared differences between model prices and market prices.

Models such as Merton's that involve jumps can be implemented with Monte Carlo simulation. When jumps are generated by a Poisson process, the probability of exactly m jumps in time t is

$$\frac{e^{-\lambda t}(\lambda t)^m}{m!}$$

where λ is the average number of jumps per year. Equivalently, λt is the average number of jumps in time t .

Suppose that on average 0.5 jumps happen per year. The probability of m jumps in 2 years is

$$\frac{e^{-0.5 \times 2}(0.5 \times 2)^m}{m!} = 0.3679$$

Table 26.1 gives the probability and cumulative probability of 0, 1, 2, 3, 4, 5, 6, 7, and 8 jumps in 2 years. (The numbers in a table such as this can be calculated using the POISSON function in Excel.)

To simulate a process following jumps over 2 years, it is necessary to determine on each simulation trial:

1. The number of jumps
2. The size of each jump.

To determine the number of jumps, on each simulation trial we sample a random number between 0 and 1 and use Table 26.1 as a look-up table. If the random number is between 0 and 0.3679, no jumps occur; if the random number is between 0.3679 and 0.7358, one jump occurs; if the random number is between 0.7358 and 0.9197, two jumps

Table 26.1 Probabilities for number of jumps in 2 years.

<i>Number of jumps, m</i>	<i>Probability of exactly m jumps</i>	<i>Probability of m jumps or less</i>
0	0.3679	0.3679
1	0.3679	0.7358
2	0.1839	0.9197
3	0.0613	0.9810
4	0.0153	0.9963
5	0.0031	0.9994
6	0.0005	0.9999
7	0.0001	1.0000
8	0.0000	1.0000

occur; and so on. To determine the size of each jump, it is necessary on each simulation trial to sample from the probability distribution for the jump size once for each jump that occurs. Once the number of jumps and the jump sizes have been determined, the final value of the variable being simulated is known for the simulation trial.

In Merton's mixed jump–diffusion model, jumps are superimposed upon the usual lognormal diffusion process that is assumed for stock prices. The process then has two components (the usual diffusion component and the jump component) and each must be sampled separately. The diffusion component is sampled as described in Sections 20.6 and 20.7 while the jump component is sampled as just described. When derivatives are valued, it is important to ensure that the overall expected return from the asset (from both components) is the risk-free rate. This means that the drift for the diffusion component in Merton's model is $r - q - \lambda k$.

The Variance-Gamma Model

An example of a pure jump model that is proving quite popular is the *variance-gamma model*.⁵ Define a variable g as the change over time T in a variable that follows a gamma process with mean rate of 1 and variance rate of v . A gamma process is a pure jump process where small jumps occur very frequently and large jumps occur only occasionally. The probability density for g is

$$\frac{g^{T/v-1} e^{-g/v}}{v^{T/v} \Gamma(T/v)}$$

where $\Gamma(\cdot)$ denotes the gamma function. This probability density can be computed in Excel using the `GAMMADIST`($\cdot, \cdot, \cdot, \cdot$) function. The first argument of the function is g , the second is T/v , the third is v , and the fourth is `TRUE` or `FALSE`, where `TRUE` returns the cumulative probability distribution function and `FALSE` returns the probability density function we have just given.

As usual, we define S_T as the asset price at time T , S_0 as the asset price today, r as the risk-free interest rate, and q as the dividend yield. In a risk-neutral world $\ln S_T$, under the variance-gamma model, has a probability distribution that, conditional on g , is normal. The conditional mean is

$$\ln S_0 + (r - q)T + \omega + \theta g$$

and the conditional standard deviation is

$$\sigma \sqrt{g}$$

where

$$\omega = (T/v) \ln(1 - \theta v - \sigma^2 v/2)$$

The variance-gamma model has three parameters: v , σ , and θ .⁶ The parameter v is the variance rate of the gamma process, σ is the volatility, and θ is a parameter defining skewness. When $\theta = 0$, $\ln S_T$ is symmetric; when $\theta < 0$, it is negatively skewed (as for equities); and when $\theta > 0$, it is positively skewed.

⁵ See D. B. Madan, P. P. Carr, and E. C. Chang, "The Variance-Gamma Process and Option Pricing," *European Finance Review*, 2 (1998): 79–105.

⁶ Note that all these parameters are liable to change when we move from the real world to the risk-neutral world. This is in contrast to pure diffusion models where the volatility remains the same.

Suppose that we are interested in using Excel to obtain 10,000 random samples of the change in an asset price between time 0 and time T using the variance-gamma model. As a preliminary, we could set cells E1, E2, E3, E4, E5, E6, and E7 equal to T , v , θ , σ , r , q , and S_0 , respectively. We could also set E8 equal to ω by defining it as

$$= \$E\$1 * LN(1 - \$E\$3 * \$E\$2 - \$E\$4 * \$E\$4 * \$E\$2/2)/\$E\$2$$

We could then proceed as follows:

1. Sample values for g using the GAMMAINV function. Set the contents of cells A1, A2, ..., A10000 as

$$= GAMMAINV(RAND(), \$E\$1/\$E\$2, \$E\$2)$$

2. For each value of g we sample a value z for a variable that is normally distributed with mean θg and standard deviation $\sigma\sqrt{g}$. This can be done by defining cell B1 as

$$= A1 * \$E\$3 + SQRT(A1) * \$E\$4 * NORMSINV(RAND())$$

and cells B2, B3, ..., B10000 similarly.

3. The stock price S_T is given by

$$S_T = S_0 \exp[(r - q)T + \omega + z]$$

By defining C1 as

$$= \$E\$7 * EXP((\$E\$5 - \$E\$6) * \$E\$1 + B1 + \$E\$8)$$

and C2, C3, ..., C10000 similarly, random samples from the distribution of S_T are created in these cells.

Figure 26.1 Distributions obtained with variance-gamma process and geometric Brownian motion.

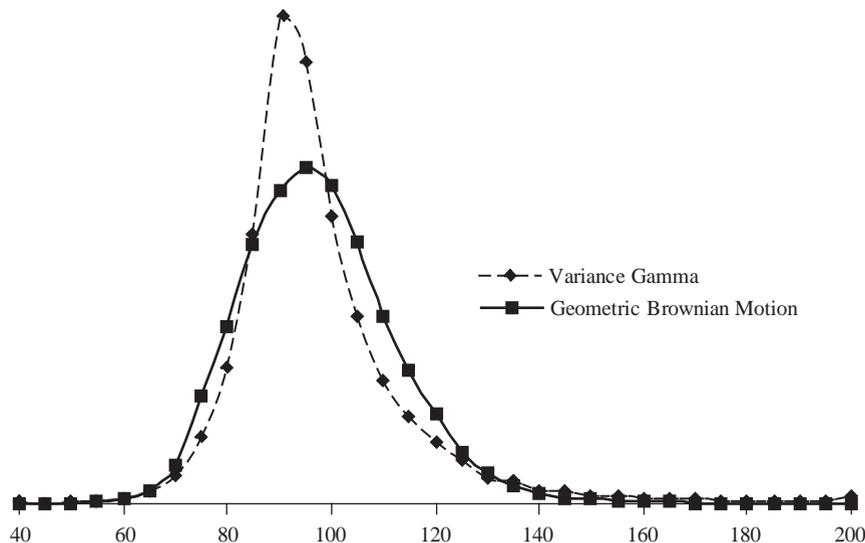


Figure 26.1 shows the probability distribution that is obtained using the variance-gamma model for S_T when $S_0 = 100$, $T = 0.5$, $v = 0.5$, $\theta = 0.1$, $\sigma = 0.2$, and $r = q = 0$. For comparison it also shows the distribution given by geometric Brownian motion when the volatility, σ is 0.2 (or 20%). Although not clear in Figure 26.1, the variance-gamma distribution has heavier tails than the lognormal distribution given by geometric Brownian motion.

One way of characterizing the variance-gamma distribution is that g defines the rate at which information arrives during time T . If g is large, a great deal of information arrives and the sample we take from a normal distribution in step 2 above has a relatively large mean and variance. If g is small, relatively little information arrives and the sample we take has a relatively small mean and variance. The parameter T is the usual time measure, and g is sometimes referred to as measuring economic time or time adjusted for the flow of information.

Semi-analytic European option valuation formulas are provided by Madan *et al.* (1998). The variance-gamma model tends to produce a U-shaped volatility smile. The smile is not necessarily symmetrical. It is very pronounced for short maturities and “dies away” for long maturities. The model can be fitted to either equity or foreign currency plain vanilla option prices.

26.2 STOCHASTIC VOLATILITY MODELS

The Black–Scholes–Merton model assumes that volatility is constant. In practice, as discussed in Chapter 22, volatility varies through time. The variance-gamma model reflects this with its g parameter. Low values of g correspond to a low arrival rate for information and a low volatility; high values of g correspond to a high arrival rate for information and a high volatility.

An alternative to the variance-gamma model is a model where the process followed by the volatility variable is specified explicitly. Suppose first that the volatility parameter in the geometric Brownian motion is a known function of time. The risk-neutral process followed by the asset price is then

$$dS = (r - q)S dt + \sigma(t)S dz \quad (26.1)$$

The Black–Scholes–Merton formulas are then correct provided that the variance rate is set equal to the average variance rate during the life of the option (see Problem 26.6). The variance rate is the square of the volatility. Suppose that during a 1-year period the volatility of a stock will be 20% during the first 6 months and 30% during the second 6 months. The average variance rate is

$$0.5 \times 0.20^2 + 0.5 \times 0.30^2 = 0.065$$

It is correct to use Black–Scholes–Merton with a variance rate of 0.065. This corresponds to a volatility of $\sqrt{0.065} = 0.255$, or 25.5%.

Equation (26.1) assumes that the instantaneous volatility of an asset is perfectly predictable. In practice, volatility varies stochastically. This has led to the development of more complex models with two stochastic variables: the stock price and its volatility.

One model that has been used by researchers is

$$\frac{dS}{S} = (r - q)dt + \sqrt{V} dz_S \quad (26.2)$$

$$dV = a(V_L - V)dt + \xi V^\alpha dz_V \quad (26.3)$$

where a , V_L , ξ , and α are constants, and dz_S and dz_V are Wiener processes. The variable V in this model is the asset's variance rate. The variance rate has a drift that pulls it back to a level V_L at rate a .

Hull and White show that, when volatility is stochastic but uncorrelated with the asset price, the price of a European option is the Black–Scholes–Merton price integrated over the probability distribution of the average variance rate during the life of the option.⁷ Thus, a European call price is

$$\int_0^\infty c(\bar{V})g(\bar{V})d\bar{V}$$

where \bar{V} is the average value of the variance rate, c is the Black–Scholes–Merton price expressed as a function of \bar{V} , and g is the probability density function of \bar{V} in a risk-neutral world. This result can be used to show that Black–Scholes–Merton overprices options that are at the money or close to the money, and underprices options that are deep in or deep out of the money. The model is consistent with the pattern of implied volatilities observed for currency options (see Section 19.2).

The case where the asset price and volatility are correlated is more complicated. Option prices can be obtained using Monte Carlo simulation. In the particular case where $\alpha = 0.5$, Hull and White provide a series expansion and Heston provides an analytic result.⁸ The pattern of implied volatilities obtained when the volatility is negatively correlated with the asset price is similar to that observed for equities (see Section 19.3).⁹

Chapter 22 discusses exponentially weighted moving average (EWMA) and GARCH(1,1) models. These are alternative approaches to characterizing a stochastic volatility model. Duan shows that it is possible to use GARCH(1,1) as the basis for an internally consistent option pricing model.¹⁰ (See Problem 22.14 for the equivalence of GARCH(1,1) and stochastic volatility models.)

Stochastic volatility models can be fitted to the prices of plain vanilla options and then used to price exotic options.¹¹ For options that last less than a year, the impact of a stochastic volatility on pricing is fairly small in absolute terms (although in percentage

⁷ See J. C. Hull and A. White, "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance*, 42 (June 1987): 281–300. This result is independent of the process followed by the variance rate.

⁸ See J. C. Hull and A. White, "An Analysis of the Bias in Option Pricing Caused by a Stochastic Volatility," *Advances in Futures and Options Research*, 3 (1988): 27–61; S. L. Heston, "A Closed Form Solution for Options with Stochastic Volatility with Applications to Bonds and Currency Options," *Review of Financial Studies*, 6, 2 (1993): 327–43.

⁹ The reason is given in footnote 3.

¹⁰ See J.-C. Duan, "The GARCH Option Pricing Model," *Mathematical Finance*, vol. 5 (1995), 13–32; and J.-C. Duan, "Cracking the Smile" *RISK*, vol. 9 (December 1996), 55–59.

¹¹ For an example of this, see J. C. Hull and W. Suo, "A Methodology for the Assessment of Model Risk and its Application to the Implied Volatility Function Model," *Journal of Financial and Quantitative Analysis*, 37, 2 (June 2002): 297–318.

terms it can be quite large for deep-out-of-the-money options). It becomes progressively larger as the life of the option increases. The impact of a stochastic volatility on the performance of delta hedging is generally quite large. Traders recognize this and, as described in Chapter 18, monitor their exposure to volatility changes by calculating vega.

26.3 THE IVF MODEL

The parameters of the models we have discussed so far can be chosen so that they provide an approximate fit to the prices of plain vanilla options on any given day. Financial institutions sometimes want to go one stage further and use a model that provides an exact fit to the prices of these options.¹² In 1994 Derman and Kani, Dupire, and Rubinstein developed a model that is designed to do this. It has become known as the *implied volatility function* (IVF) model or the *implied tree* model.¹³ It provides an exact fit to the European option prices observed on any given day, regardless of the shape of the volatility surface.

The risk-neutral process for the asset price in the model has the form

$$dS = [r(t) - q(t)]S dt + \sigma(S, t)S dz$$

where $r(t)$ is the instantaneous forward interest rate for a contract maturing at time t and $q(t)$ is the dividend yield as a function of time. The volatility $\sigma(S, t)$ is a function of both S and t and is chosen so that the model prices all European options consistently with the market. It is shown both by Dupire and by Andersen and Brotherton-Ratcliffe that $\sigma(S, t)$ can be calculated analytically:¹⁴

$$[\sigma(K, T)]^2 = 2 \frac{\partial c_{\text{mkt}} / \partial T + q(T)c_{\text{mkt}} + K[r(T) - q(T)]\partial c_{\text{mkt}} / \partial K}{K^2(\partial^2 c_{\text{mkt}} / \partial K^2)} \quad (26.4)$$

where $c_{\text{mkt}}(K, T)$ is the market price of a European call option with strike price K and maturity T . If a sufficiently large number of European call prices are available in the market, this equation can be used to estimate the $\sigma(S, t)$ function.¹⁵

Andersen and Brotherton-Ratcliffe implement the model by using equation (26.4) together with the implicit finite difference method. An alternative approach, the *implied tree* methodology suggested by Derman and Kani and Rubinstein, involves constructing a tree for the asset price that is consistent with option prices in the market.

When it is used in practice the IVF model is recalibrated daily to the prices of plain vanilla options. It is a tool to price exotic options consistently with plain vanilla options. As discussed in Chapter 19 plain vanilla options define the risk-neutral

¹² There is a practical reason for this. If the bank does not use a model with this property, there is a danger that traders working for the bank will spend their time arbitraging the bank's internal models.

¹³ See B. Dupire, "Pricing with a Smile," *Risk*, February (1994): 18–20; E. Derman and I. Kani, "Riding on a Smile," *Risk*, February (1994): 32–39; M. Rubinstein, "Implied Binomial Trees" *Journal of Finance*, 49, 3 (July 1994), 771–818.

¹⁴ See B. Dupire, "Pricing with a Smile," *Risk*, February (1994), 18–20; L.B.G. Andersen and R. Brotherton-Ratcliffe "The Equity Option Volatility Smile: An Implicit Finite Difference Approach," *Journal of Computation Finance* 1, No. 2 (Winter 1997/98): 5–37. Dupire considers the case where r and q are zero; Andersen and Brotherton-Ratcliffe consider the more general situation.

¹⁵ Some smoothing of the observed volatility surface is typically necessary.

probability distribution of the asset price at all future times. It follows that the IVF model gets the risk-neutral probability distribution of the asset price at all future times correct. This means that options providing payoffs at just one time (e.g., all-or-nothing and asset-or-nothing options) are priced correctly by the IVF model. However, the model does not necessarily get the joint distribution of the asset price at two or more times correct. This means that exotic options such as compound options and barrier options may be priced incorrectly.¹⁶

26.4 CONVERTIBLE BONDS

We now move on to discuss how the numerical procedures presented in Chapter 20 can be modified to handle particular valuation problems. We start by considering convertible bonds.

Convertible bonds are bonds issued by a company where the holder has the option to exchange the bonds for the company's stock at certain times in the future. The *conversion ratio* is the number of shares of stock obtained for one bond (this can be a function of time). The bonds are almost always callable (i.e., the issuer has the right to buy them back at certain times at a predetermined prices). The holder always has the right to convert the bond once it has been called. The call feature is therefore usually a way of forcing conversion earlier than the holder would otherwise choose. Sometimes the holder's call option is conditional on the price of the company's stock being above a certain level.

Credit risk plays an important role in the valuation of convertibles. If credit risk is ignored, poor prices are obtained because the coupons and principal payments on the bond are overvalued. Ingersoll provides a way of valuing convertibles using a model similar to Merton's (1974) model discussed in Section 23.6.¹⁷ He assumes geometric Brownian motion for the issuer's total assets and models the company's equity, its convertible debt, and its other debt as claims contingent on the value of the assets. Credit risk is taken into account because the debt holders get repaid in full only if the value of the assets exceeds the amount owing to them.

A simpler model that is widely used in practice involves modeling the issuer's stock price. It is assumed that the stock follows geometric Brownian motion except that there is a probability $\lambda \Delta t$ that there will be a default in each short period of time Δt . In the event of a default the stock price falls to zero and there is a recovery on the bond. The variable λ is the risk-neutral default intensity defined in Section 23.2.

The stock price process can be represented by varying the usual binomial tree so that at each node there is:

1. A probability p_u of a percentage up movement of size u over the next time period of length Δt

¹⁶ Hull and Suo test the IVF model by assuming that all derivative prices are determined by a stochastic volatility model. They found that the model works reasonably well for compound options, but sometimes gives serious errors for barrier options. See J.C. Hull and W. Suo, "A Methodology for the Assessment of Model Risk and its Application to the Implied Volatility Function Model," *Journal of Financial and Quantitative Analysis*, 37, 2 (June 2002): 297–318

¹⁷ See J.E. Ingersoll, "A Contingent Claims Valuation of Convertible Securities," *Journal of Financial Economics*, 4, (May 1977), 289–322.

2. A probability p_d of a percentage down movement of size d over the next time period of length Δt
3. A probability $\lambda \Delta t$, or more accurately $1 - e^{-\lambda \Delta t}$, that there will be a default with the stock price moving to zero over the next time period of length Δt

Parameter values, chosen to match the first two moments of the stock price distribution, are:

$$p_u = \frac{a - de^{-\lambda \Delta t}}{u - d}, \quad p_d = \frac{ue^{-\lambda \Delta t} - a}{u - d}, \quad u = e^{\sqrt{(\sigma^2 - \lambda)\Delta t}}, \quad d = \frac{1}{u}$$

where $a = e^{(r-q)\Delta t}$, r is the risk-free rate, and q is the dividend yield on the stock.

The life of the tree is set equal to the life of the convertible bond. The value of the convertible at the final nodes of the tree is calculated based on any conversion options that the holder has at that time. We then roll back through the tree. At nodes where the terms of the instrument allow conversion we test whether conversion is optimal. We also test whether the position of the issuer can be improved by calling the bonds. If so, we assume that the bonds are called and retest whether conversion is optimal. This is equivalent to setting the value at a node equal to

$$\max[\min(Q_1, Q_2), Q_3]$$

where Q_1 is the value given by the rollback (assuming that the bond is neither converted nor called at the node), Q_2 is the call price, and Q_3 is the value if conversion takes place.

Example 26.1

Consider a 9-month zero-coupon bond issued by company XYZ with a face value of \$100. Suppose that it can be exchanged for two shares of company XYZ's stock at any time during the 9 months. Assume also that it is callable for \$113 at any time. The initial stock price is \$50, its volatility is 30% per annum, and there are no dividends. The default intensity λ is 1% per year, and all risk-free rates for all maturities are 5%. Suppose that in the event of a default the bond is worth \$40 (i.e., the recovery rate, as it is usually defined, is 40%).

Figure 26.2 shows the stock price tree that can be used to value the convertible when there are three time steps ($\Delta t = 0.25$). The upper number at each node is the stock price; the lower number is the price of the convertible bond. The tree parameters are:

$$u = e^{\sqrt{(0.09 - 0.01) \times 0.25}} = 1.1519, \quad d = 1/u = 0.8681$$

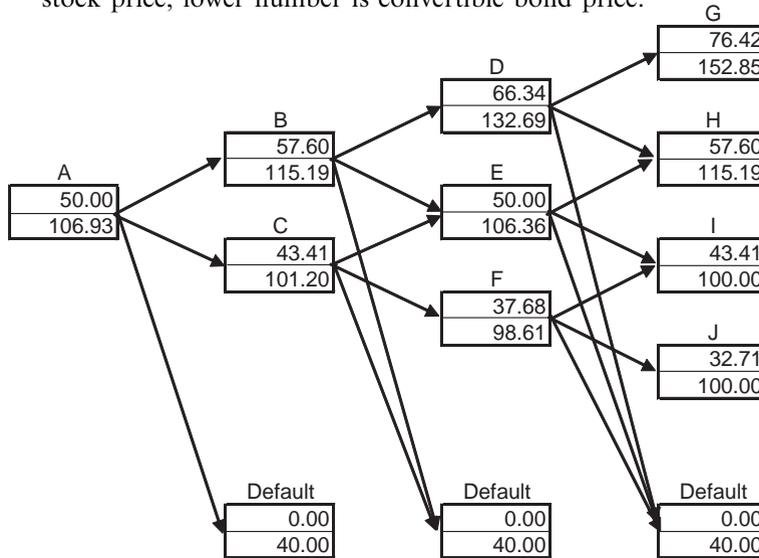
$$a = e^{0.05 \times 0.25} = 1.0126, \quad p_u = 0.5167, \quad p_d = 0.4808$$

The probability of a default (i.e., of moving to the lowest nodes on the tree is $1 - e^{-0.01 \times 0.25} = 0.002497$. At the three default nodes the stock price is zero and the bond price is 40.

Consider first the final nodes. At nodes G and H the bond should be converted and is worth twice the stock price. At nodes I and J the bond should not be converted and is worth 100.

Moving back through the tree enables the value to be calculated at earlier nodes. Consider, for example, node E. The value, if the bond is converted, is $2 \times 50 = \$100$. If it is not converted, then there is (a) a probability 0.5167 that it will move to node H, where the bond is worth 115.19, (b) a 0.4808 probability

Figure 26.2 Tree for valuing convertible. Upper number at each node is stock price; lower number is convertible bond price.



that it will move down to node I, where the bond is worth 100, and (c) a 0.002497 probability that it will default and be worth 40. The value of the bond if it is not converted is therefore

$$(0.5167 \times 115.19 + 0.4808 \times 100 + 0.002497 \times 40) \times e^{-0.05 \times 0.25} = 106.36$$

This is more than the value of 100 that it would have if converted. We deduce that it is not worth converting the bond at node E. Finally, we note that the bond issuer would not call the bond at node E because this would be offering 113 for a bond worth 106.36.

As another example consider node B. The value of the bond if it is converted is $2 \times 57.596 = 115.19$. If it is not converted a similar calculation to that just given for node E gives its value as 118.31. The convertible bond holder will therefore choose not to convert. However, at this stage the bond issuer will call the bond for 113 and the bond holder will then decide that converting is better than being called. The value of the bond at node B is therefore 115.19. A similar argument is used to arrive at the value at node D. With no conversion the value is 132.79. However, the bond is called, forcing conversion and reducing the value at the node to 132.69.

The value of the convertible is its value at the initial node A, or 106.93.

When interest is paid on the debt, it must be taken into account. At each node, when valuing the bond on the assumption that it is not converted, the present value of any interest payable on the bond in the next time step should be included. The risk-neutral default intensity λ can be estimated from either bond prices or credit default swap spreads. In a more general implementation, λ , σ , and r are functions of time. This can be handled using a trinomial rather than a binomial tree (see Section 20.4).

One disadvantage of the model we have presented is that the probability of default is independent of the stock price. This has led some researchers to suggest an implicit finite difference method implementation of the model where the default intensity λ is a function of the stock price as well as time.¹⁸

26.5 PATH-DEPENDENT DERIVATIVES

A path-dependent derivative (or history-dependent derivative) is a derivative where the payoff depends on the path followed by the price of the underlying asset, not just its final value. Asian options and lookback options are examples of path-dependent derivatives. As explained in Chapter 25, the payoff from an Asian option depends on the average price of the underlying asset; the payoff from a lookback option depends on its maximum or minimum price. One approach to valuing path-dependent options when analytic results are not available is Monte Carlo simulation, as discussed in Chapter 20. A sample value of the derivative can be calculated by sampling a random path for the underlying asset in a risk-neutral world, calculating the payoff, and discounting the payoff at the risk-free interest rate. An estimate of the value of the derivative is found by obtaining many sample values of the derivative in this way and calculating their mean.

The main problem with Monte Carlo simulation is that the computation time necessary to achieve the required level of accuracy can be unacceptably high. Also, American-style path-dependent derivatives (i.e., path-dependent derivatives where one side has exercise opportunities or other decisions to make) cannot easily be handled. In this section, we show how the binomial tree methods presented in Chapter 20 can be extended to cope with some path-dependent derivatives.¹⁹ The procedure can handle American-style path-dependent derivatives and is computationally more efficient than Monte Carlo simulation for European-style path-dependent derivatives.

For the procedure to work, two conditions must be satisfied:

1. The payoff from the derivative must depend on a single function, F , of the path followed by the underlying asset.
2. It must be possible to calculate the value of F at time $\tau + \Delta t$ from the value of F at time τ and the value of the underlying asset at time $\tau + \Delta t$.

Illustration Using Lookback Options

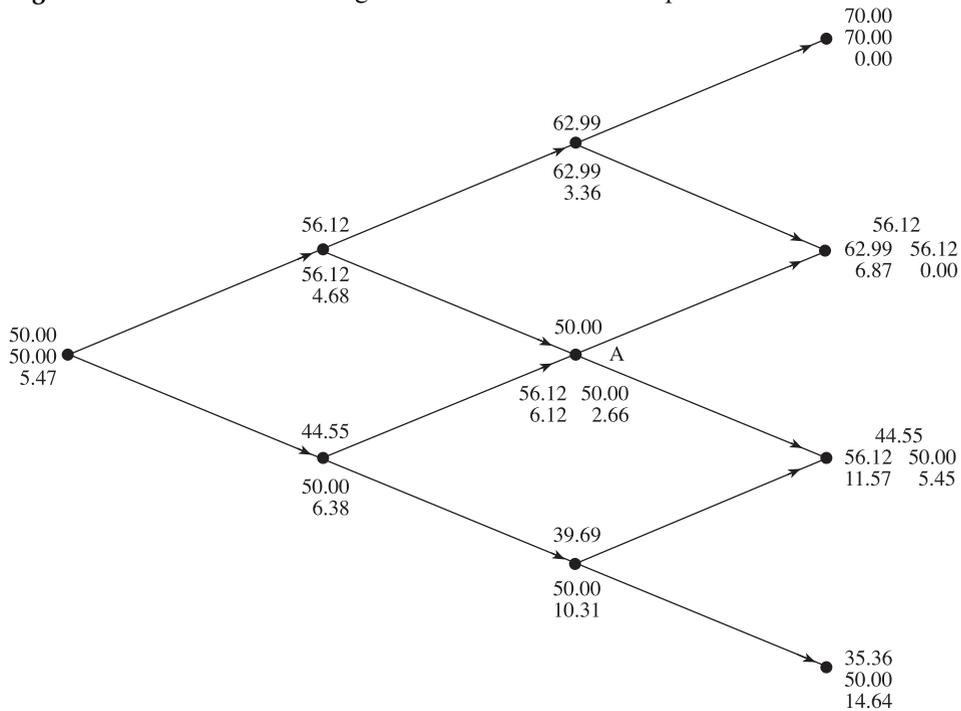
As a first illustration of the procedure, consider an American floating lookback put option on a non-dividend-paying stock.²⁰ If exercised at time τ , this pays off the amount by which the maximum stock price between time 0 and time τ exceeds the current stock

¹⁸ See, e.g., L. Andersen and D. Buffum, "Calibration and Implementation of Convertible Bond Models," *Journal of Computational Finance*, 7, 1 (Winter 2003/04), 1–34. These authors suggest assuming that the default intensity is inversely proportional to S^α , where S is the stock price and α is a positive constant.

¹⁹ This approach was suggested in J. Hull and A. White, "Efficient Procedures for Valuing European and American Path-Dependent Options," *Journal of Derivatives*, 1, 1 (Fall 1993): 21–31.

²⁰ This example is used as a first illustration of the general procedure for handling path dependence. For a more efficient approach to valuing American-style lookback options, see Technical Note 13 at:

www.rotman.utoronto.ca/~hull/TechnicalNotes.

Figure 26.3 Tree for valuing an American lookback option.

price. Suppose that the initial stock price is \$50, the stock price volatility is 40% per annum, the risk-free interest rate is 10% per annum, the total life of the option is three months, and that stock price movements are represented by a three-step binomial tree. With our usual notation this means that $S_0 = 50$, $\sigma = 0.4$, $r = 0.10$, $\Delta t = 0.08333$, $u = 1.1224$, $d = 0.8909$, $a = 1.0084$, and $p = 0.5073$.

The tree is shown in Figure 26.3. In this case, the path function F is the maximum stock price so far. The top number at each node is the stock price. The next level of numbers at each node shows the possible maximum stock prices achievable on paths leading to the node. The final level of numbers shows the values of the derivative corresponding to each of the possible maximum stock prices.

The values of the derivative at the final nodes of the tree are calculated as the maximum stock price minus the actual stock price. To illustrate the rollback procedure, suppose that we are at node A, where the stock price is \$50. The maximum stock price achieved thus far is either 56.12 or 50. Consider first the situation where it is equal to 50. If there is an up movement, the maximum stock price becomes 56.12 and the value of the derivative is zero. If there is a down movement, the maximum stock price stays at 50 and the value of the derivative is 5.45. Assuming no early exercise, the value of the derivative at A when the maximum achieved so far is 50 is, therefore,

$$(0 \times 0.5073 + 5.45 \times 0.4927)e^{-0.1 \times 0.08333} = 2.66$$

Clearly, it is not worth exercising at node A in these circumstances because the payoff

from doing so is zero. A similar calculation for the situation where the maximum value at node A is 56.12 gives the value of the derivative at node A, without early exercise, to be

$$(0 \times 0.5073 + 11.57 \times 0.4927)e^{-0.1 \times 0.08333} = 5.65$$

In this case, early exercise gives a value of 6.12 and is the optimal strategy. Rolling back through the tree in the way we have indicated gives the value of the American lookback as \$5.47.

Generalization

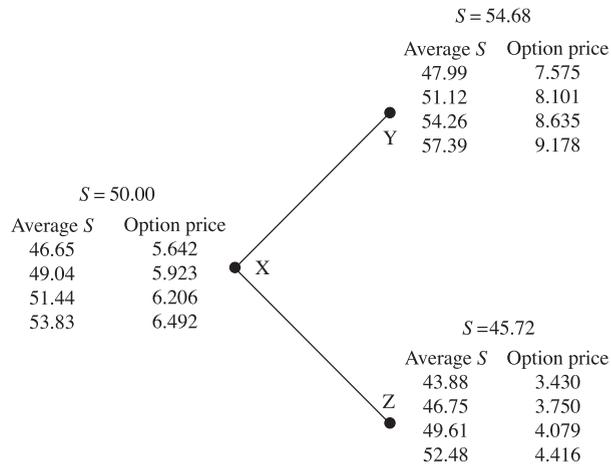
The approach just described is computationally feasible when the number of alternative values of the path function, F , at each node does not grow too fast as the number of time steps is increased. The example we used, a lookback option, presents no problems because the number of alternative values for the maximum asset price at a node in a binomial tree with n time steps is never greater than n .

Luckily, the approach can be extended to cope with situations where there are a very large number of different possible values of the path function at each node. The basic idea is as follows. Calculations are carried out at each node for a small number of representative values of F . When the value of the derivative is required for other values of the path function, it is calculated from the known values using interpolation.

The first stage is to work forward through the tree establishing the maximum and minimum values of the path function at each node. Assuming the value of the path function at time $\tau + \Delta t$ depends only on the value of the path function at time τ and the value of the underlying variable at time $\tau + \Delta t$, the maximum and minimum values of the path function for the nodes at time $\tau + \Delta t$ can be calculated in a straightforward way from those for the nodes at time τ . The second stage is to choose representative values of the path function at each node. There are a number of approaches. A simple rule is to choose the representative values as the maximum value, the minimum value, and a number of other values that are equally spaced between the maximum and the minimum. As we roll back through the tree, we value the derivative for each of the representative values of the path function.

To illustrate the nature of the calculation, consider the problem of valuing the average price call option in Example 25.2 of Section 25.12 when the payoff depends on the arithmetic average stock price. The initial stock price is 50, the strike price is 50, the risk-free interest rate is 10%, the stock price volatility is 40%, and the time to maturity is 1 year. For 20 time steps, the binomial tree parameters are $\Delta t = 0.05$, $u = 1.0936$, $d = 0.9144$, $p = 0.5056$, and $1 - p = 0.4944$. The path function is the arithmetic average of the stock price.

Figure 26.4 shows the calculations that are carried out in one small part of the tree. Node X is the central node at time 0.2 year (at the end of the fourth time step). Nodes Y and Z are the two nodes at time 0.25 year that are reachable from node X. The stock price at node X is 50. Forward induction shows that the maximum average stock price that is achievable in reaching node X is 53.83. The minimum is 46.65. (The initial and final stock prices are included when calculating the average.) From node X, the tree branches to one of the two nodes Y and Z. At node Y, the stock price is 54.68 and the bounds for the average are 47.99 and 57.39. At node Z, the stock price is 45.72 and the bounds for the average stock price are 43.88 and 52.48.

Figure 26.4 Part of tree for valuing option on the arithmetic average.

Suppose that the representative values of the average are chosen to be four equally spaced values at each node. This means that, at node X, averages of 46.65, 49.04, 51.44, and 53.83 are considered. At node Y, the averages 47.99, 51.12, 54.26, and 57.39 are considered. At node Z, the averages 43.88, 46.75, 49.61, and 52.48 are considered. Assume that backward induction has already been used to calculate the value of the option for each of the alternative values of the average at nodes Y and Z. Values are shown in Figure 26.4 (e.g., at node Y when the average is 51.12, the value of the option is 8.101).

Consider the calculations at node X for the case where the average is 51.44. If the stock price moves up to node Y, the new average will be

$$\frac{5 \times 51.44 + 54.68}{6} = 51.98$$

The value of the derivative at node Y for this average can be found by interpolating between the values when the average is 51.12 and when it is 54.26. It is

$$\frac{(51.98 - 51.12) \times 8.635 + (54.26 - 51.98) \times 8.101}{54.26 - 51.12} = 8.247$$

Similarly, if the stock price moves down to node Z, the new average will be

$$\frac{5 \times 51.44 + 45.72}{6} = 50.49$$

and by interpolation the value of the derivative is 4.182.

The value of the derivative at node X when the average is 51.44 is, therefore,

$$(0.5056 \times 8.247 + 0.4944 \times 4.182)e^{-0.1 \times 0.05} = 6.206$$

The other values at node X are calculated similarly. Once the values at all nodes at time 0.2 year have been calculated, the nodes at time 0.15 year can be considered.

The value given by the full tree for the option at time zero is 7.17. As the number of time steps and the number of averages considered at each node is increased, the value of the option converges to the correct answer. With 60 time steps and 100 averages at each node, the value of the option is 5.58. The analytic approximation for the value of the option, as calculated in Example 25.2, with continuous averaging is 5.62.

A key advantage of the method described here is that it can handle American options. The calculations are as we have described them except that we test for early exercise at each node for each of the alternative values of the path function at the node. (In practice, the early exercise decision is liable to depend on both the value of the path function and the value of the underlying asset.) Consider the American version of the average price call considered here. The value calculated using the 20-step tree and four averages at each node is 7.77; with 60 time steps and 100 averages, the value is 6.17.

The approach just described can be used in a wide range of different situations. The two conditions that must be satisfied were listed at the beginning of this section. Efficiency is improved somewhat if quadratic rather than linear interpolation is used at each node.

26.6 BARRIER OPTIONS

Chapter 25 presented analytic results for standard barrier options. This section considers numerical procedures that can be used for barrier options when there are no analytic results.

In principle, many barrier options can be valued using the binomial and trinomial trees discussed in Chapter 20. Consider an up-and-out option. A simple approach is to value this in the same way as a regular option except that, when a node above the barrier is encountered, the value of the option is set equal to zero.

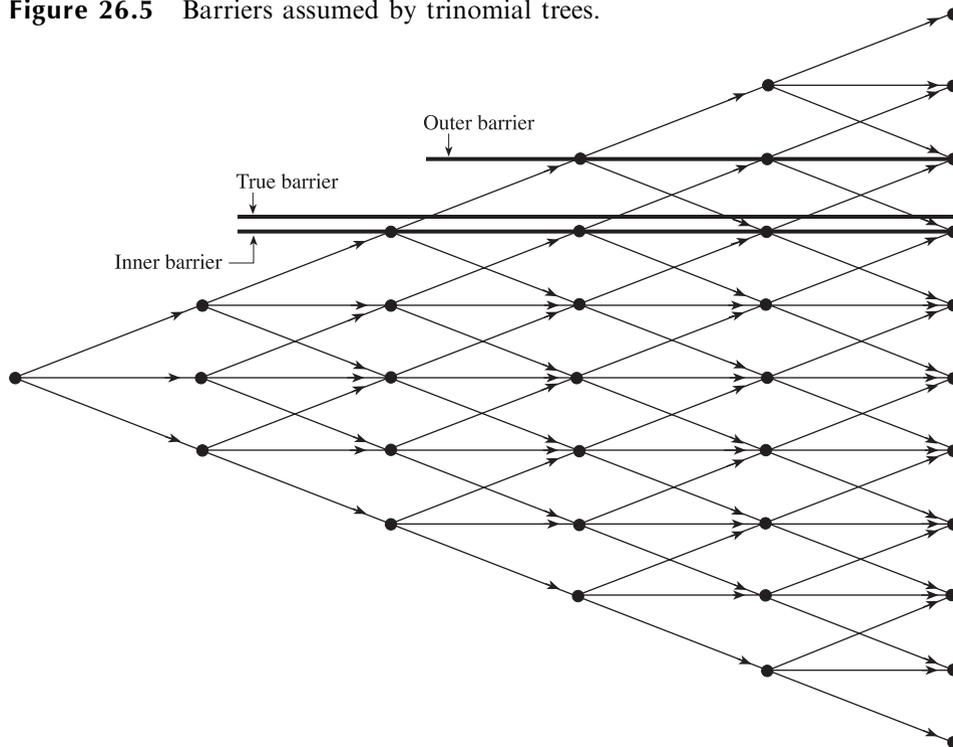
Trinomial trees work better than binomial trees, but even for them convergence is very slow when the simple approach is used. A large number of time steps are required to obtain a reasonably accurate result. The reason for this is that the barrier being assumed by the tree is different from the true barrier.²¹ Define the *inner barrier* as the barrier formed by nodes just on the inside of the true barrier (i.e., closer to the center of the tree) and the *outer barrier* as the barrier formed by nodes just outside the true barrier (i.e., farther away from the center of the tree). Figure 26.5 shows the inner and outer barrier for a trinomial tree on the assumption that the true barrier is horizontal. The usual tree calculations implicitly assume that the outer barrier is the true barrier because the barrier conditions are first used at nodes on this barrier. When the time step is Δt , the vertical spacing between the nodes is of order $\sqrt{\Delta t}$. This means that errors created by the difference between the true barrier and the outer barrier also tend to be of order $\sqrt{\Delta t}$.

One approach to overcoming this problem is to:

1. Calculate the price of the derivative on the assumption that the inner barrier is the true barrier.

²¹ For a discussion of this, see P.P. Boyle and S.H. Lau, "Bumping Up Against the Barrier with the Binomial Method," *Journal of Derivatives*, 1, 4 (Summer 1994): 6–14.

Figure 26.5 Barriers assumed by trinomial trees.



2. Calculate the value of the derivative on the assumption that the outer barrier is the true barrier.
3. Interpolate between the two prices.

Another approach is to ensure that nodes lie on the barrier. Suppose that the initial stock price is S_0 and that the barrier is at H . In a trinomial tree, there are three possible movements in the asset's price at each node: up by a proportional amount u ; stay the same; and down by a proportional amount d , where $d = 1/u$. We can always choose u so that nodes lie on the barrier. The condition that must be satisfied by u is

$$H = S_0 u^N$$

or

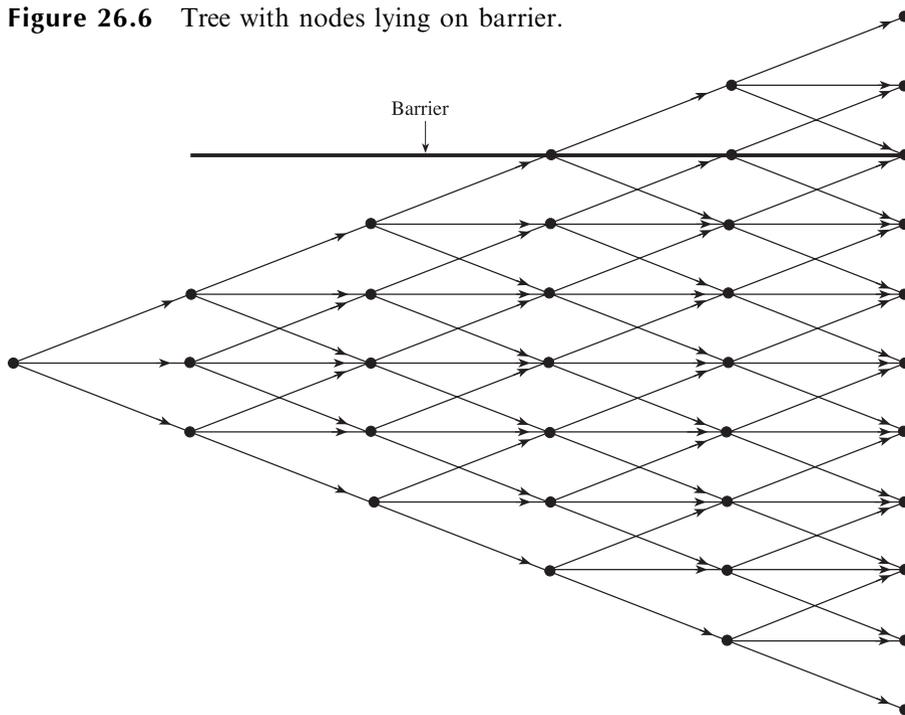
$$\ln H = \ln S_0 + N \ln u$$

for some positive or negative N .

When discussing trinomial trees in Section 20.4, the value suggested for u was $e^{\sigma\sqrt{3\Delta t}}$, so that $\ln u = \sigma\sqrt{3\Delta t}$. In the situation considered here, a good rule is to choose $\ln u$ as close as possible to this value, consistent with the condition given above. This means that

$$\ln u = \frac{\ln H - \ln S_0}{N}$$

Figure 26.6 Tree with nodes lying on barrier.



where

$$N = \text{int} \left[\frac{\ln H - \ln S_0}{\sigma \sqrt{3\Delta t}} + 0.5 \right]$$

and $\text{int}(x)$ is the integral part of x .

This leads to a tree of the form shown in Figure 26.6. The probabilities p_u , p_m , and p_d on the upper, middle, and lower branches of the tree are chosen to match the first two moments of the return, so that

$$p_d = -\frac{(r - q - \sigma^2/2)\Delta t}{2 \ln u} + \frac{\sigma^2 \Delta t}{2(\ln u)^2}, \quad p_m = 1 - \frac{\sigma^2 \Delta t}{(\ln u)^2}, \quad p_u = \frac{(r - q - \sigma^2/2)\Delta t}{2 \ln u} + \frac{\sigma^2 \Delta t}{2(\ln u)^2}$$

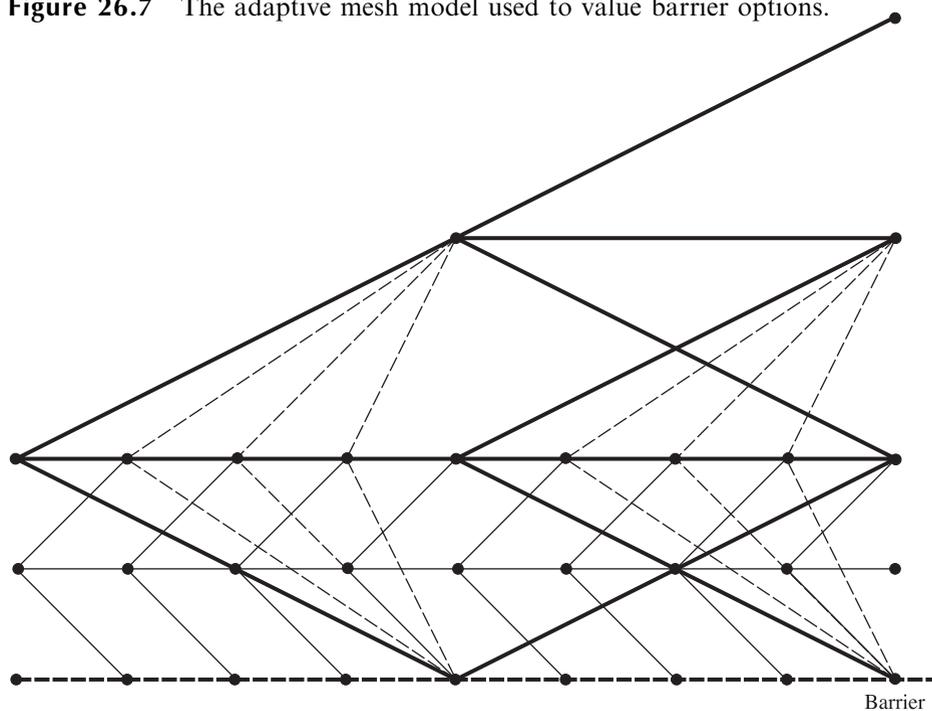
where p_u , p_m , and p_d are the probabilities on the upper, middle, and lower branches.

The Adaptive Mesh Model

The methods presented so far work reasonably well when the initial asset price is not close to the barrier. When the initial asset price is close to a barrier, the adaptive mesh model, which was introduced in Section 20.4, can be used.²² The idea behind the model is that computational efficiency can be improved by grafting a fine tree onto a coarse

²² See S. Figlewski and B. Gao, “The Adaptive Mesh Model: A New Approach to Efficient Option Pricing,” *Journal of Financial Economics*, 53 (1999): 313–51.

Figure 26.7 The adaptive mesh model used to value barrier options.



tree to achieve a more detailed modeling of the asset price in the regions of the tree where it is needed most.

To value a barrier option, it is useful to have a fine tree close to barriers. Figure 26.7 illustrates the design of the tree. The geometry of the tree is arranged so that nodes lie on the barriers. The probabilities on branches are chosen, as usual, to match the first two moments of the process followed by the underlying asset. The heavy lines in Figure 26.7 are the branches of the coarse tree. The light solid lines are the fine tree. We first roll back through the coarse tree in the usual way. We then calculate the value at additional nodes using the branches indicated by the dotted lines. Finally we roll back through the fine tree.

26.7 OPTIONS ON TWO CORRELATED ASSETS

Another tricky numerical problem is that of valuing American options dependent on two assets whose prices are correlated. A number of alternative approaches have been suggested. This section will explain three of these.

Transforming Variables

It is relatively easy to construct a tree in three dimensions to represent the movements of two *uncorrelated* variables. The procedure is as follows. First, construct a two-dimensional tree for each variable, and then combine these trees into a single three-

dimensional tree. The probabilities on the branches of the three-dimensional tree are the product of the corresponding probabilities on the two-dimensional trees. Suppose, for example, that the variables are stock prices, S_1 and S_2 . Each can be represented in two dimensions by a Cox, Ross, and Rubinstein binomial tree. Assume that S_1 has a probability p_1 of moving up by a proportional amount u_1 and a probability $1 - p_1$ of moving down by a proportional amount d_1 . Suppose further that S_2 has a probability p_2 of moving up by a proportional amount u_2 and a probability $1 - p_2$ of moving down by a proportional amount d_2 . In the three-dimensional tree there are four branches emanating from each node. The probabilities are:

$$\begin{aligned} p_1 p_2: & S_1 \text{ increases; } S_2 \text{ increases} \\ p_1(1 - p_2): & S_1 \text{ increases; } S_2 \text{ decreases} \\ (1 - p_1)p_2: & S_1 \text{ decreases; } S_2 \text{ increases} \\ (1 - p_1)(1 - p_2): & S_1 \text{ decreases; } S_2 \text{ decreases} \end{aligned}$$

Consider next the situation where S_1 and S_2 are correlated. Suppose that the risk-neutral processes are:

$$\begin{aligned} dS_1 &= (r - q_1)S_1 dt + \sigma_1 S_1 dz_1 \\ dS_2 &= (r - q_2)S_2 dt + \sigma_2 S_2 dz_2 \end{aligned}$$

and the instantaneous correlation between the Wiener processes, dz_1 and dz_2 , is ρ . This means that

$$\begin{aligned} d \ln S_1 &= (r - q_1 - \sigma_1^2/2) dt + \sigma_1 dz_1 \\ d \ln S_2 &= (r - q_2 - \sigma_2^2/2) dt + \sigma_2 dz_2 \end{aligned}$$

Two new uncorrelated variables can be defined:²³

$$\begin{aligned} x_1 &= \sigma_2 \ln S_1 + \sigma_1 \ln S_2 \\ x_2 &= \sigma_2 \ln S_1 - \sigma_1 \ln S_2 \end{aligned}$$

These variables follow the processes

$$\begin{aligned} dx_1 &= [\sigma_2(r - q_1 - \sigma_1^2/2) + \sigma_1(r - q_2 - \sigma_2^2/2)] dt + \sigma_1 \sigma_2 \sqrt{2(1 + \rho)} dz_A \\ dx_2 &= [\sigma_2(r - q_1 - \sigma_1^2/2) - \sigma_1(r - q_2 - \sigma_2^2/2)] dt + \sigma_1 \sigma_2 \sqrt{2(1 - \rho)} dz_B \end{aligned}$$

where dz_A and dz_B are uncorrelated Wiener processes.

The variables x_1 and x_2 can be modeled using two separate binomial trees. In time Δt , x_i has a probability p_i of increasing by h_i and a probability $1 - p_i$ of decreasing by h_i . The variables h_i and p_i are chosen so that the tree gives correct values for the first two moments of the distribution of x_1 and x_2 . Because they are uncorrelated, the two trees can be combined into a single three-dimensional tree, as already described. At each node of the tree, S_1 and S_2 can be calculated from x_1 and x_2 using the inverse

²³ This idea was suggested in J. Hull and A. White, "Valuing Derivative Securities Using the Explicit Finite Difference Method," *Journal of Financial and Quantitative Analysis*, 25 (1990): 87–100.

relationships

$$S_1 = \exp\left[\frac{x_1 + x_2}{2\sigma_2}\right] \quad \text{and} \quad S_2 = \exp\left[\frac{x_1 - x_2}{2\sigma_1}\right]$$

The procedure for rolling back through a three-dimensional tree to value a derivative is analogous to that for a two-dimensional tree.

Using a Nonrectangular Tree

Rubinstein has suggested a way of building a three-dimensional tree for two correlated stock prices by using a nonrectangular arrangement of the nodes.²⁴ From a node (S_1, S_2) , where the first stock price is S_1 and the second stock price is S_2 , there is a 0.25 chance of moving to each of the following:

$$(S_1u_1, S_2A), \quad (S_1u_1, S_2B), \quad (S_1d_1, S_2C), \quad (S_1d_1, S_2D)$$

where

$$u_1 = \exp[(r - q_1 - \sigma_1^2/2)\Delta t + \sigma_1\sqrt{\Delta t}]$$

$$d_1 = \exp[(r - q_1 - \sigma_1^2/2)\Delta t - \sigma_1\sqrt{\Delta t}]$$

and

$$A = \exp[(r - q_2 - \sigma_2^2/2)\Delta t + \sigma_2\sqrt{\Delta t}(\rho + \sqrt{1 - \rho^2})]$$

$$B = \exp[(r - q_2 - \sigma_2^2/2)\Delta t + \sigma_2\sqrt{\Delta t}(\rho - \sqrt{1 - \rho^2})]$$

$$C = \exp[(r - q_2 - \sigma_2^2/2)\Delta t - \sigma_2\sqrt{\Delta t}(\rho - \sqrt{1 - \rho^2})]$$

$$D = \exp[(r - q_2 - \sigma_2^2/2)\Delta t - \sigma_2\sqrt{\Delta t}(\rho + \sqrt{1 - \rho^2})]$$

When the correlation is zero, this method is equivalent to constructing separate trees for S_1 and S_2 using the alternative binomial tree construction method in Section 19.4.

Adjusting the Probabilities

A third approach to building a three-dimensional tree for S_1 and S_2 involves first assuming no correlation and then adjusting the probabilities at each node to reflect the correlation.²⁵ The alternative binomial tree construction method for each of S_1 and S_2 in Section 20.4 is used. This method has the property that all probabilities are 0.5. When the

Table 26.2 Combination of binomials assuming no correlation.

S_2 -move	S_1 -move	
	Down	Up
Up	0.25	0.25
Down	0.25	0.25

²⁴ See M. Rubinstein, "Return to Oz," *Risk*, November (1994): 67–70.

²⁵ This approach was suggested in the context of interest rate trees in J. Hull and A. White, "Numerical Procedures for Implementing Term Structure Models II: Two-Factor Models," *Journal of Derivatives*, Winter (1994): 37–48.

Table 26.3 Combination of binomials assuming correlation of ρ .

S_2 -move	S_1 -move	
	Down	Up
Up	$0.25(1 - \rho)$	$0.25(1 + \rho)$
Down	$0.25(1 + \rho)$	$0.25(1 - \rho)$

two binomial trees are combined on the assumption that there is no correlation, the probabilities are as shown in Table 26.2. When the probabilities are adjusted to reflect the correlation, they become those shown in Table 26.3.

26.8 MONTE CARLO SIMULATION AND AMERICAN OPTIONS

Monte Carlo simulation is well suited to valuing path-dependent options and options where there are many stochastic variables. Trees and finite difference methods are well suited to valuing American-style options. What happens if an option is both path dependent and American? What happens if an American option depends on several stochastic variables? Section 26.5 explained a way in which the binomial tree approach can be modified to value path-dependent options in some situations. A number of researchers have adopted a different approach by searching for a way in which Monte Carlo simulation can be used to value American-style options.²⁶ This section explains two alternative ways of proceeding.

The Least-Squares Approach

In order to value an American-style option it is necessary to choose between exercising and continuing at each early exercise point. The value of exercising is normally easy to determine. A number of researchers including Longstaff and Schwartz provide a way of determining the value of continuing when Monte Carlo simulation is used.²⁷ Their approach involves using a least-squares analysis to determine the best-fit relationship between the value of continuing and the values of relevant variables at each time an early exercise decision has to be made. The approach is best illustrated with a numerical example. We use the one in the Longstaff–Schwartz paper.

Consider a 3-year American put option on a non-dividend-paying stock that can be exercised at the end of year 1, the end of year 2, and the end of year 3. The risk-free rate is 6% per annum (continuously compounded). The current stock price is 1.00 and the strike price is 1.10. Assume that the eight paths shown in Table 26.4 are sampled for the stock price. (This example is for illustration only; in practice many more paths would be sampled.) If the option can be exercised only at the 3-year point, it provides a cash flow equal to its intrinsic value at that point. This is shown in the last column of Table 26.5.

²⁶ Tilley was the first researcher to publish a solution to the problem. See J. A. Tilley, “Valuing American Options in a Path Simulation Model,” *Transactions of the Society of Actuaries*, 45 (1993): 83–104.

²⁷ See F. A. Longstaff and E. S. Schwartz, “Valuing American Options by Simulation: A Simple Least-Squares Approach,” *Review of Financial Studies*, 14, 1 (Spring 2001): 113–47.

Table 26.4 Sample paths for put option example.

<i>Path</i>	$t = 0$	$t = 1$	$t = 2$	$t = 3$
1	1.00	1.09	1.08	1.34
2	1.00	1.16	1.26	1.54
3	1.00	1.22	1.07	1.03
4	1.00	0.93	0.97	0.92
5	1.00	1.11	1.56	1.52
6	1.00	0.76	0.77	0.90
7	1.00	0.92	0.84	1.01
8	1.00	0.88	1.22	1.34

If the put option is in the money at the 2-year point, the option holder must decide whether to exercise. Table 26.4 shows that the option is in the money at the 2-year point for paths 1, 3, 4, 6, and 7. For these paths, we assume an approximate relationship:

$$V = a + bS + cS^2$$

where S is the stock price at the 2-year point and V is the value of continuing, discounted back to the 2-year point. Our five observations on S are: 1.08, 1.07, 0.97, 0.77, and 0.84. From Table 26.5 the corresponding values for V are: 0.00 , $0.07e^{-0.06 \times 1}$, $0.18e^{-0.06 \times 1}$, $0.20e^{-0.06 \times 1}$, and $0.09e^{-0.06 \times 1}$. The values of a , b , and c that minimize

$$\sum_{i=1}^5 (V_i - a - bS_i - cS_i^2)^2$$

where S_i and V_i are the i th observation on S and V , respectively, are $a = -1.070$, $b = 2.983$ and $c = -1.813$, so that the best-fit relationship is

$$V = -1.070 + 2.983S - 1.813S^2$$

This gives the value at the 2-year point of continuing for paths 1, 3, 4, 6, and 7 of 0.0369, 0.0461, 0.1176, 0.1520, and 0.1565, respectively. From Table 26.4 the value of exercising

Table 26.5 Cash flows if exercise only possible at 3-year point.

<i>Path</i>	$t = 1$	$t = 2$	$t = 3$
1	0.00	0.00	0.00
2	0.00	0.00	0.00
3	0.00	0.00	0.07
4	0.00	0.00	0.18
5	0.00	0.00	0.00
6	0.00	0.00	0.20
7	0.00	0.00	0.09
8	0.00	0.00	0.00

Table 26.6 Cash flows if exercise only possible at 2- and 3-year point.

<i>Path</i>	$t = 1$	$t = 2$	$t = 3$
1	0.00	0.00	0.00
2	0.00	0.00	0.00
3	0.00	0.00	0.07
4	0.00	0.13	0.00
5	0.00	0.00	0.00
6	0.00	0.33	0.00
7	0.00	0.26	0.00
8	0.00	0.00	0.00

is 0.02, 0.03, 0.13, 0.33, and 0.26. This means that we should exercise at the 2-year point for paths 4, 6, and 7. Table 26.6 summarizes the cash flows assuming exercise at either the 2-year point or the 3-year point for the eight paths.

Consider next the paths that are in the money at the 1-year point. These are paths 1, 4, 6, 7, and 8. From Table 26.4 the values of S for the paths are 1.09, 0.93, 0.76, 0.92, and 0.88, respectively. From Table 26.6, the corresponding continuation values discounted back to $t = 1$ are 0.00 , $0.13e^{-0.06 \times 1}$, $0.33e^{-0.06 \times 1}$, $0.26e^{-0.06 \times 1}$, and 0.00 , respectively. The least-squares relationship is

$$V = 2.038 - 3.335S + 1.356S^2$$

This gives the value of continuing at the 1-year point for paths 1, 4, 6, 7, 8 as 0.0139, 0.1092, 0.2866, 0.1175, and 0.1533, respectively. From Table 26.4 the value of exercising is 0.01, 0.17, 0.34, 0.18, and 0.22, respectively. This means that we should exercise at the 1-year point for paths 4, 6, 7, and 8. Table 26.7 summarizes the cash flows assuming that early exercise is possible at all three times. The value of the option is determined by discounting each cash flow back to time zero at the risk-free rate and calculating the mean of the results. It is

$$\frac{1}{8}(0.07e^{-0.06 \times 3} + 0.17e^{-0.06 \times 1} + 0.34e^{-0.06 \times 1} + 0.18e^{-0.06 \times 1} + 0.22e^{-0.06 \times 1}) = 0.1144$$

Since this is greater than 0.10, it is not optimal to exercise the option immediately.

Table 26.7 Cash flows from option.

<i>Path</i>	$t = 1$	$t = 2$	$t = 3$
1	0.00	0.00	0.00
2	0.00	0.00	0.00
3	0.00	0.00	0.07
4	0.17	0.00	0.00
5	0.00	0.00	0.00
6	0.34	0.00	0.00
7	0.18	0.00	0.00
8	0.22	0.00	0.00

This method can be extended in a number of ways. If the option can be exercised at any time we can approximate its value by considering a large number of exercise points (just as a binomial tree does). The relationship between V and S can be assumed to be more complicated. For example we could assume that V is a cubic rather than a quadratic function of S . The method can be used where the early exercise decision depends on several state variables. A functional form for the relationship between V and the variables is assumed and the parameters are estimated using the least-squares approach, as in the example just considered.

The Exercise Boundary Parameterization Approach

A number of researchers, such as Andersen, have proposed an alternative approach where the early exercise boundary is parameterized and the optimal values of the parameters are determined iteratively by starting at the end of the life of the option and working backward.²⁸ To illustrate the approach, we continue with the put option example and assume that the eight paths shown in Table 26.4 have been sampled. In this case, the early exercise boundary at time t can be parameterized by a critical value of S , $S^*(t)$. If the asset price at time t is below $S^*(t)$ we exercise at time t ; if it is above $S^*(t)$ we do not exercise at time t . The value of $S^*(3)$ is 1.10. If the stock price is above 1.10 when $t = 3$ (the end of the option's life) we do not exercise; if it is below 1.10 we exercise. We now consider the determination of $S^*(2)$.

Suppose that we choose a value of $S^*(2)$ less than 0.77. The option is not exercised at the 2-year point for any of the paths. The value of the option at the 2-year point for the eight paths is then 0.00, 0.00, $0.07e^{-0.06 \times 1}$, $0.18e^{-0.06 \times 1}$, 0.00, $0.20e^{-0.06 \times 1}$, $0.09e^{-0.06 \times 1}$, and 0.00, respectively. The average of these is 0.0636. Suppose next that $S^*(2) = 0.77$. The value of the option at the 2-year point for the eight paths is then 0.00, 0.00, $0.07e^{-0.06 \times 1}$, $0.18e^{-0.06 \times 1}$, 0.00, 0.33, $0.09e^{-0.06 \times 1}$, and 0.00, respectively. The average of these is 0.0813. Similarly when $S^*(2)$ equals 0.84, 0.97, 1.07, and 1.08, the average value of the option at the 2-year point is 0.1032, 0.0982, 0.0938, and 0.0963, respectively. This analysis shows that the optimal value of $S^*(2)$ (i.e., the one that maximizes the average value of the option) is 0.84. (More precisely, it is optimal to choose $0.84 \leq S^*(2) < 0.97$.) When we choose this optimal value for $S^*(2)$, the value of the option at the 2-year point for the eight paths is 0.00, 0.00, 0.0659, 0.1695, 0.00, 0.33, 0.26, and 0.00, respectively. The average value is 0.1032.

We now move on to calculate $S^*(1)$. If $S^*(1) < 0.76$ the option is not exercised at the 1-year point for any of the paths and the value at the option at the 1-year point is $0.1032e^{-0.06 \times 1} = 0.0972$. If $S^*(1) = 0.76$, the value of the option for each of the eight paths at the 1-year point is 0.00, 0.00, $0.0659e^{-0.06 \times 1}$, $0.1695e^{-0.06 \times 1}$, 0.0, 0.34, $0.26e^{-0.06 \times 1}$, and 0.00, respectively. The average value of the option is 0.1008. Similarly when $S^*(1)$ equals 0.88, 0.92, 0.93, and 1.09 the average value of the option is 0.1283, 0.1202, 0.1215, and 0.1228, respectively. The analysis therefore shows that the optimal value of $S^*(1)$ is 0.88. (More precisely, it is optimal to choose $0.88 \leq S^*(1) < 0.92$.) The value of the option at time zero with no early exercise is $0.1283e^{-0.06 \times 1} = 0.1208$. This is greater than the value of 0.10 obtained by exercising at time zero.

In practice, tens of thousands of simulations are carried out to determine the early exercise boundary in the way we have described. Once the early exercise boundary has

²⁸ See L. Andersen, "A Simple Approach to the Pricing of Bermudan Swaptions in the Multifactor LIBOR Market Model," *Journal of Computational Finance*, 3, 2 (Winter 2000): 1–32.

been obtained, the paths for the variables are discarded and a new Monte Carlo simulation using the early exercise boundary is carried out to value the option. Our American put option example is simple in that we know that the early exercise boundary at a time can be defined entirely in terms of the value of the stock price at that time. In more complicated situations it is necessary to make assumptions about how the early exercise boundary should be parameterized.

Upper Bounds

The two approaches we have outlined tend to underprice American-style options because they assume a suboptimal early exercise boundary. This has led Andersen and Broadie to propose a procedure that provides an upper bound to the price.²⁹ This procedure can be used in conjunction with any algorithm that generates a lower bound and pinpoints the true value of an American-style option more precisely than the algorithm does by itself.

SUMMARY

A number of models have been developed to fit the volatility smiles that are observed in practice. The constant elasticity of variance model leads to a volatility smile similar to that observed for equity options. The jump–diffusion model leads to a volatility smile similar to that observed for currency options. Variance-gamma and stochastic volatility models are more flexible in that they can lead to either the type of volatility smile observed for equity options or the type of volatility smile observed for currency options. The implied volatility function model provides even more flexibility than this. It is designed to provide an exact fit to any pattern of European option prices observed in the market.

The natural technique to use for valuing path-dependent options is Monte Carlo simulation. This has the disadvantage that it is fairly slow and unable to handle American-style derivatives easily. Luckily, trees can be used to value many types of path-dependent derivatives. The approach is to choose representative values for the underlying path function at each node of the tree and calculate the value of the derivative for each of these values as we roll back through the tree.

The binomial tree methodology can be extended to value convertible bonds. Extra branches corresponding to a default by the company are added to the tree. The roll-back calculations then reflect the holder's option to convert and the issuer's option to call.

Trees can be used to value many types of barrier options, but the convergence of the option value to the correct value as the number of time steps is increased tends to be slow. One approach for improving convergence is to arrange the geometry of the tree so that nodes always lie on the barriers. Another is to use an interpolation scheme to adjust for the fact that the barrier being assumed by the tree is different from the true barrier. A third is to design the tree so that it provides a finer representation of movements in the underlying asset price near the barrier.

One way of valuing options dependent on the prices of two correlated assets is to apply a transformation to the asset price to create two new uncorrelated variables.

²⁹ See L. Andersen and M. Broadie, "A Primal-Dual Simulation Algorithm for Pricing Multi-Dimensional American Options," *Management Science*, 50, 9 (2004), 1222–34.

These two variables are each modeled with trees and the trees are then combined to form a single three-dimensional tree. At each node of the tree, the inverse of the transformation gives the asset prices. A second approach is to arrange the positions of nodes on the three-dimensional tree to reflect the correlation. A third approach is to start with a tree that assumes no correlation between the variables and then adjust the probabilities on the tree to reflect the correlation.

Monte Carlo simulation is not naturally suited to valuing American-style options, but there are two ways it can be adapted to handle them. The first uses a least-squares analysis to relate the value of continuing (i.e., not exercising) to the values of relevant variables. The second involves parameterizing the early exercise boundary and determining it iteratively by working back from the end of the life of the option to the beginning.

FURTHER READING

- Andersen, L., "A Simple Approach to the Pricing of Bermudan Swaptions in the Multifactor LIBOR Market Model," *Journal of Computational Finance*, 3, 2 (Winter 2000): 1–32.
- Andersen, L. B. G., and R. Brotherton-Ratcliffe, "The Equity Option Volatility Smile: An Implicit Finite Difference Approach," *Journal of Computational Finance*, 1, 2 (Winter 1997/98): 3–37.
- Bodurtha, J. N., and M. Jermakyan, "Non-Parametric Estimation of an Implied Volatility Surface," *Journal of Computational Finance*, 2, 4 (Summer 1999): 29–61.
- Boyle, P. P., and S. H. Lau, "Bumping Up Against the Barrier with the Binomial Method," *Journal of Derivatives*, 1, 4 (Summer 1994): 6–14.
- Conze, A., and R. Viswanathan, "Path Dependent Options: The Case of Lookback Options," *Journal of Finance*, 46 (1991): 1893–1907.
- Cox, J. C. and S. A. Ross, "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics*, 3 (March 1976), 145–66.
- Derman, E., and I. Kani, "Riding on a Smile," *Risk*, February (1994): 32–39.
- Duan, J.-C., "The GARCH Option Pricing Model," *Mathematical Finance*, 5 (1995): 13–32.
- Duan, J.-C., "Cracking the Smile," *Risk*, December (1996): 55–59.
- Dupire, B., "Pricing with a Smile," *Risk*, February (1994): 18–20.
- Figlewski, S., and B. Gao, "The Adaptive Mesh Model: A New Approach to Efficient Option Pricing," *Journal of Financial Economics*, 53 (1999): 313–51.
- Heston, S. L., "A Closed Form Solution for Options with Stochastic Volatility with Applications to Bonds and Currency Options," *Review of Financial Studies*, 6, 2 (1993): 327–43.
- Hull, J., and A. White, "Efficient Procedures for Valuing European and American Path-Dependent Options," *Journal of Derivatives*, 1, 1 (Fall 1993): 21–31.
- Hull J. C., and A. White, "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance*, 42 (June 1987): 281–300.
- Hull, J. C. and W. Suo, "A Methodology for the Assessment of Model Risk and its Application to the Implied Volatility Function Model," *Journal of Financial and Quantitative Analysis*, 37, 2 (2002): 297–318.
- Longstaff, F. A. and E. S. Schwartz, "Valuing American Options by Simulation: A Simple Least-Squares Approach," *Review of Financial Studies*, 14, 1 (Spring 2001): 113–47.
- Madan D. B., P. P. Carr, and E. C. Chang, "The Variance-Gamma Process and Option Pricing" *European Finance Review*, 2 (1998): 7–105.
- Merton, R. C., "Option Pricing When Underlying Stock Returns Are Discontinuous," *Journal of Financial Economics*, 3 (March 1976): 125–44.

- Rebonato, R., *Volatility and Correlation: The Perfect Hedger and the Fox*, 2nd edn. Chichester: Wiley, 2004.
- Ritchken, P, and R. Trevor, “Pricing Options Under Generalized GARCH and Stochastic Volatility Processes,” *Journal of Finance* 54, 1 (February 1999): 377–402
- Rubinstein, M., “Implied Binomial Trees,” *Journal of Finance*, 49, 3 (July 1994): 771–818.
- Rubinstein, M., “Return to Oz,” *Risk*, November (1994): 67–70.
- Stutzer, M., “A Simple Nonparametric Approach to Derivative Security Valuation,” *Journal of Finance*, 51 (December 1996): 1633–52.
- Tilley, J.A., “Valuing American Options in a Path Simulation Model,” *Transactions of the Society of Actuaries*, 45 (1993): 83–104.

Practice Questions (Answers in Solutions Manual)

- 26.1. Confirm that the CEV model formulas satisfy put–call parity.
- 26.2. Use Monte Carlo simulation to show that Merton’s value for a European option is correct when $r = 0.05$, $q = 0$, $\lambda = 0.3$, $k = 0.5$, $\sigma = 0.25$, and $S_0 = 30$.
- 26.3. Confirm that Merton’s jump–diffusion model satisfies put–call parity when the jump size is lognormal.
- 26.4. Suppose that the volatility of an asset will be 20% from month 0 to month 6, 22% from month 6 to month 12, and 24% from month 12 to month 24. What volatility should be used in Black–Scholes–Merton to value a 2-year option?
- 26.5. Consider the case of Merton’s jump–diffusion model where jumps always reduce the asset price to zero. Assume that the average number of jumps per year is λ . Show that the price of a European call option is the same as in a world with no jumps except that the risk-free rate is $r + \lambda$ rather than r . Does the possibility of jumps increase or reduce the value of the call option in this case? (*Hint*: Value the option assuming no jumps and assuming one or more jumps. The probability of no jumps in time T is $e^{-\lambda T}$).
- 26.6. At time 0 the price of a non-dividend-paying stock is S_0 . Suppose that the time interval between 0 and T is divided into two subintervals of length t_1 and t_2 . During the first subinterval, the risk-free interest rate and volatility are r_1 and σ_1 , respectively. During the second subinterval, they are r_2 and σ_2 , respectively. Assume that the world is risk neutral.
- Use the results in Chapter 14 to determine the stock price distribution at time T in terms of r_1 , r_2 , σ_1 , σ_2 , t_1 , t_2 , and S_0 .
 - Suppose that \bar{r} is the average interest rate between time zero and T and that \bar{V} is the average variance rate between times zero and T . What is the stock price distribution as a function of T in terms of \bar{r} , \bar{V} , T , and S_0 ?
 - What are the results corresponding to (a) and (b) when there are three subintervals with different interest rates and volatilities?
 - Show that if the risk-free rate, r , and the volatility, σ , are known functions of time, the stock price distribution at time T in a risk-neutral world is

$$\ln S_T \sim \phi[\ln S_0 + (\bar{r} - \frac{1}{2}\bar{V})T, VT]$$

where \bar{r} is the average value of r , \bar{V} is equal to the average value of σ^2 , and S_0 is the stock price today and $\phi(m, v)$ is a normal distribution with mean m and variance v .

- 26.7. Write down the equations for simulating the path followed by the asset price in the stochastic volatility model in equations (26.2) and (26.3).
- 26.8. “The IVF model does not necessarily get the evolution of the volatility surface correct.” Explain this statement.
- 26.9. “When interest rates are constant the IVF model correctly values any derivative whose payoff depends on the value of the underlying asset at only one time.” Explain why.
- 26.10. Use a three-time-step tree to value an American floating lookback call option on a currency when the initial exchange rate is 1.6, the domestic risk-free rate is 5% per annum, the foreign risk-free interest rate is 8% per annum, the exchange rate volatility is 15%, and the time to maturity is 18 months. Use the approach in Section 26.5.
- 26.11. What happens to the variance-gamma model as the parameter v tends to zero?
- 26.12. Use a three-time-step tree to value an American put option on the geometric average of the price of a non-dividend-paying stock when the stock price is \$40, the strike price is \$40, the risk-free interest rate is 10% per annum, the volatility is 35% per annum, and the time to maturity is three months. The geometric average is measured from today until the option matures.
- 26.13. Can the approach for valuing path-dependent options in Section 26.5 be used for a 2-year American-style option that provides a payoff equal to $\max(S_{\text{ave}} - K, 0)$, where S_{ave} is the average asset price over the three months preceding exercise? Explain your answer.
- 26.14. Verify that the 6.492 number in Figure 26.4 is correct.
- 26.15. Examine the early exercise policy for the eight paths considered in the example in Section 26.8. What is the difference between the early exercise policy given by the least squares approach and the exercise boundary parameterization approach? Which gives a higher option price for the paths sampled?
- 26.16. Consider a European put option on a non-dividend paying stock when the stock price is \$100, the strike price is \$110, the risk-free rate is 5% per annum, and the time to maturity is one year. Suppose that the average variance rate during the life of an option has a 0.20 probability of being 0.06, a 0.5 probability of being 0.09, and a 0.3 probability of being 0.12. The volatility is uncorrelated with the stock price. Estimate the value of the option. Use DerivaGem.
- 26.17. When there are two barriers how can a tree be designed so that nodes lie on both barriers?
- 26.18. Consider an 18-month zero-coupon bond with a face value of \$100 that can be converted into five shares of the company’s stock at any time during its life. Suppose that the current share price is \$20, no dividends are paid on the stock, the risk-free rate for all maturities is 6% per annum with continuous compounding, and the share price volatility is 25% per annum. Assume that the default intensity is 3% per year and the recovery rate is 35%. The bond is callable at \$110. Use a three-time-step tree to calculate the value of the bond. What is the value of the conversion option (net of the issuer’s call option)?

Further Questions

- 26.19. A new European-style floating lookback call option on a stock index has a maturity of 9 months. The current level of the index is 400, the risk-free rate is 6% per annum, the dividend yield on the index is 4% per annum, and the volatility of the index is 20%. Use

the approach in Section 26.5 to value the option and compare your answer to the result given by DerivaGem using the analytic valuation formula.

- 26.20. Suppose that the volatilities used to price a 6-month currency option are as in Table 19.2. Assume that the domestic and foreign risk-free rates are 5% per annum and the current exchange rate is 1.00. Consider a bull spread that consists of a long position in a 6-month call option with strike price 1.05 and a short position in a 6-month call option with a strike price 1.10.
- What is the value of the spread?
 - What single volatility if used for both options gives the correct value of the bull spread? (Use the DerivaGem Application Builder in conjunction with Goal Seek or Solver.)
 - Does your answer support the assertion at the beginning of the chapter that the correct volatility to use when pricing exotic options can be counterintuitive?
 - Does the IVF model give the correct price for the bull spread?
- 26.21. Repeat the analysis in Section 26.8 for the put option example on the assumption that the strike price is 1.13. Use both the least squares approach and the exercise boundary parameterization approach.
- 26.22. Consider the situation in Merton's jump-diffusion model where the underlying asset is a non-dividend-paying stock. The average frequency of jumps is one per year. The average percentage jump size is 2% and the standard deviation of the logarithm of the percentage jump size is 20%. The stock price is 100, the risk-free rate is 5%, the volatility, σ provided by the diffusion part of the process is 15%, and the time to maturity is six months. Use the DerivaGem Application Builder to calculate an implied volatility when the strike price is 80, 90, 100, 110, and 120. What does the volatility smile or skew that you obtain imply about the probability distribution of the stock price.
- 26.23. A 3-year convertible bond with a face value of \$100 has been issued by company ABC. It pays a coupon of \$5 at the end of each year. It can be converted into ABC's equity at the end of the first year or at the end of the second year. At the end of the first year, it can be exchanged for 3.6 shares immediately after the coupon date. At the end of the second year, it can be exchanged for 3.5 shares immediately after the coupon date. The current stock price is \$25 and the stock price volatility is 25%. No dividends are paid on the stock. The risk-free interest rate is 5% with continuous compounding. The yield on bonds issued by ABC is 7% with continuous compounding and the recovery rate is 30%.
- Use a three-step tree to calculate the value of the bond.
 - How much is the conversion option worth?
 - What difference does it make to the value of the bond and the value of the conversion option if the bond is callable any time within the first 2 years for \$115?
 - Explain how your analysis would change if there were a dividend payment of \$1 on the equity at the 6-month, 18-month, and 30-month points. Detailed calculations are not required.
- (Hint: Use equation (23.2) to estimate the average default intensity.)

Chapter 29 uses it to value some nonstandard derivatives, and Chapter 31 uses it to develop the LIBOR market model.

27.1 THE MARKET PRICE OF RISK

We start by considering the properties of derivatives dependent on the value of a single variable θ . Assume that the process followed by θ is

$$\frac{d\theta}{\theta} = m dt + s dz \quad (27.1)$$

where dz is a Wiener process. The parameters m and s are the expected growth rate in θ and the volatility of θ , respectively. We assume that they depend only on θ and time t . The variable θ need not be the price of an investment asset. It could be something as far removed from financial markets as the temperature in the center of New Orleans.

Suppose that f_1 and f_2 are the prices of two derivatives dependent only on θ and t . These can be options or other instruments that provide a payoff equal to some function of θ at some future time. Assume that during the time period under consideration f_1 and f_2 provide no income.¹

Suppose that the processes followed by f_1 and f_2 are

$$\frac{df_1}{f_1} = \mu_1 dt + \sigma_1 dz$$

and

$$\frac{df_2}{f_2} = \mu_2 dt + \sigma_2 dz$$

where μ_1 , μ_2 , σ_1 , and σ_2 are functions of θ and t . The “ dz ” in these processes must be the same dz as in equation (27.1) because it is the only source of the uncertainty in the prices of f_1 and f_2 .

The prices f_1 and f_2 can be related using an analysis similar to the Black–Scholes analysis described in Section 14.6. The discrete versions of the processes for f_1 and f_2 are

$$\Delta f_1 = \mu_1 f_1 \Delta t + \sigma_1 f_1 \Delta z \quad (27.2)$$

$$\Delta f_2 = \mu_2 f_2 \Delta t + \sigma_2 f_2 \Delta z \quad (27.3)$$

We can eliminate the Δz by forming an instantaneously riskless portfolio consisting of $\sigma_2 f_2$ of the first derivative and $-\sigma_1 f_1$ of the second derivative. If Π is the value of the portfolio, then

$$\Pi = (\sigma_2 f_2) f_1 - (\sigma_1 f_1) f_2 \quad (27.4)$$

and

$$\Delta \Pi = \sigma_2 f_2 \Delta f_1 - \sigma_1 f_1 \Delta f_2$$

Substituting from equations (27.2) and (27.3), this becomes

$$\Delta \Pi = (\mu_1 \sigma_2 f_1 f_2 - \mu_2 \sigma_1 f_1 f_2) \Delta t \quad (27.5)$$

¹ The analysis can be extended to derivatives that provide income (see Problem 27.7).

Because the portfolio is instantaneously riskless, it must earn the risk-free rate. Hence,

$$\Delta\Pi = r\Pi \Delta t$$

Substituting into this equation from equations (27.4) and (27.5) gives

$$\mu_1\sigma_2 - \mu_2\sigma_1 = r\sigma_2 - r\sigma_1$$

or

$$\frac{\mu_1 - r}{\sigma_1} = \frac{\mu_2 - r}{\sigma_2} \quad (27.6)$$

Note that the left-hand side of equation (27.6) depends only on the parameters of the process followed by f_1 and the right-hand side depends only on the parameters of the process followed by f_2 . Define λ as the value of each side in equation (27.6), so that

$$\frac{\mu_1 - r}{\sigma_1} = \frac{\mu_2 - r}{\sigma_2} = \lambda$$

Dropping subscripts, equation (27.6) shows that if f is the price of a derivative dependent only on θ and t with

$$\frac{df}{f} = \mu dt + \sigma dz \quad (27.7)$$

then

$$\frac{\mu - r}{\sigma} = \lambda \quad (27.8)$$

The parameter λ is known as the *market price of risk* of θ . (In the context of portfolio performance measurement, it is known as the *Sharpe ratio*.) It can be dependent on both θ and t , but it is not dependent on the nature of the derivative f . Our analysis shows that, for no arbitrage, $(\mu - r)/\sigma$ must at any given time be the same for all derivatives that are dependent only on θ and t .

The market price of risk of θ measures the trade-offs between risk and return that are made for securities dependent on θ . Equation (27.8) can be written

$$\mu - r = \lambda\sigma \quad (27.9)$$

The variable σ can be loosely interpreted as the quantity of θ -risk present in f . On the right-hand side of the equation, the quantity of θ -risk is multiplied by the price of θ -risk. The left-hand side is the expected return, in excess of the risk-free interest rate, that is required to compensate for this risk. Equation (27.9) is analogous to the capital asset pricing model, which relates the expected excess return on a stock to its risk. This chapter will not be concerned with the measurement of the market price of risk. This will be discussed in Chapter 34 when the evaluation of real options is considered.

It is natural to assume that σ , the coefficient of dz , in equation (27.8) is the volatility of f . In fact, σ can be negative. This will be the case when f is negatively related to θ (so that $\partial f/\partial\theta$ is negative). It is the absolute value $|\sigma|$ of σ that is the volatility of f . One way of understanding this is to note that the process for f has the same statistical properties when we replace dz by $-dz$.

Chapter 5 distinguished between investment assets and consumption assets. An investment asset is an asset that is bought or sold purely for investment purposes by a significant number of investors. Consumption assets are held primarily for consumption.

Equation (27.8) is true for all investment assets that provide no income and depend only on θ . If the variable θ itself happens to be such an asset, then

$$\frac{m - r}{s} = \lambda$$

But, in other circumstances, this relationship is not necessarily true.

Example 27.1

Consider a derivative whose price is positively related to the price of oil and depends on no other stochastic variables. Suppose that it provides an expected return of 12% per annum and has a volatility of 20% per annum. Assume that the risk-free interest rate is 8% per annum. It follows that the market price of risk of oil is

$$\frac{0.12 - 0.08}{0.2} = 0.2$$

Note that oil is a consumption asset rather than an investment asset, so its market price of risk cannot be calculated from equation (27.8) by setting μ equal to the expected return from an investment in oil and σ equal to the volatility of oil prices.

Example 27.2

Consider two securities, both of which are positively dependent on the 90-day interest rate. Suppose that the first one has an expected return of 3% per annum and a volatility of 20% per annum, and the second one has a volatility of 30% per annum. Assume that the instantaneous risk-free rate of interest is 6% per annum. The market price of interest rate risk is, using the expected return and volatility for the first security,

$$\frac{0.03 - 0.06}{0.2} = -0.15$$

From a rearrangement of equation (27.9), the expected return from the second security is, therefore,

$$0.06 - 0.15 \times 0.3 = 0.015$$

or 1.5% per annum.

Alternative Worlds

The process followed by derivative price f is

$$df = \mu f dt + \sigma f dz$$

The value of μ depends on the risk preferences of investors. In a world where the market price of risk is zero, λ equals zero. From equation (27.9) $\mu = r$, so that the process followed by f is

$$df = rf dt + \sigma f dz$$

We will refer to this as the *traditional risk-neutral world*.

Other assumptions about the market price of risk, λ , enable other worlds that are internally consistent to be defined. From equation (27.9),

$$\mu = r + \lambda\sigma$$

so that

$$df = (r + \lambda\sigma)f dt + \sigma f dz \quad (27.10)$$

The market price of risk of a variable determines the growth rates of all securities dependent on the variable. As we move from one market price of risk to another, the expected growth rates of security prices change, but their volatilities remain the same. This is a general property of variables following diffusion processes and was illustrated in Section 12.7. Choosing a particular market price of risk is also referred to as defining the *probability measure*. Some value of the market price of risk corresponds to the “real world” and the growth rates of security prices that are observed in practice.

27.2 SEVERAL STATE VARIABLES

Suppose that n variables, $\theta_1, \theta_2, \dots, \theta_n$, follow stochastic processes of the form

$$d\theta_i/\theta_i = m_i dt + s_i dz_i \quad (27.11)$$

for $i = 1, 2, \dots, n$, where the dz_i are Wiener processes. The parameters m_i and s_i are expected growth rates and volatilities and may be functions of the θ_i and time. Equation (13A.10) in the appendix to Chapter 13 provides a version of Itô’s lemma that covers functions of several variables. It shows that the process for the price f of a security that is dependent on the θ_i has n stochastic components. It can be written

$$df/f = \mu dt + \sum_{i=1}^n \sigma_i dz_i \quad (27.12)$$

In this equation, μ is the expected return from the security and $\sigma_i dz_i$ is the component of the risk of this return attributable to θ_i . Both μ and the σ_i are potentially dependent on the θ_i and time.

Technical Note 30 at www.rotman.utoronto.ca/~hull/TechnicalNotes shows that

$$\mu - r = \sum_{i=1}^n \lambda_i \sigma_i \quad (27.13)$$

where λ_i is the market price of risk for θ_i . This equation relates the expected excess return that investors require on the security to the λ_i and σ_i . Equation (27.9) is the particular case of this equation when $n = 1$. The term $\lambda_i \sigma_i$ on the right-hand side measures the extent that the excess return required by investors on a security is affected by the dependence of the security on θ_i . If $\lambda_i \sigma_i = 0$, there is no effect; if $\lambda_i \sigma_i > 0$, investors require a higher return to compensate them for the risk arising from θ_i ; if $\lambda_i \sigma_i < 0$, the dependence of the security on θ_i causes investors to require a lower return than would otherwise be the case. The $\lambda_i \sigma_i < 0$ situation occurs when the variable has the effect of reducing rather than increasing the risks in the portfolio of a typical investor.

Example 27.3

A stock price depends on three underlying variables: the price of oil, the price of gold, and the performance of a stock index. Suppose that the market prices of risk for these variables are 0.2, -0.1 , and 0.4, respectively. Suppose also that the σ_i in

equation (27.12) corresponding to the three variables have been estimated as 0.05, 0.1, and 0.15, respectively. The excess return on the stock over the risk-free rate is

$$0.2 \times 0.05 - 0.1 \times 0.1 + 0.4 \times 0.15 = 0.06$$

or 6.0% per annum. If variables other than those considered affect the stock price, this result is still true provided that the market price of risk for each of these other variables is zero.

Equation (27.13) is closely related to arbitrage pricing theory, developed by Stephen Ross in 1976.² The continuous-time version of the capital asset pricing model (CAPM) can be regarded as a particular case of the equation. CAPM argues that an investor requires excess returns to compensate for any risk that is correlated to the risk in the return from the stock market, but requires no excess return for other risks. Risks that are correlated with the return from the stock market are referred to as *systematic*; other risks are referred to as *nonsystematic*. If CAPM is true, then λ_i is proportional to the correlation between changes in θ_i and the return from the market. When θ_i is uncorrelated with the return from the market, λ_i is zero.

27.3 MARTINGALES

A *martingale* is a zero-drift stochastic process.³ A variable θ follows a martingale if its process has the form

$$d\theta = \sigma dz$$

where dz is a Wiener process. The variable σ may itself be stochastic. It can depend on θ and other stochastic variables. A martingale has the convenient property that its expected value at any future time is equal to its value today. This means that

$$E(\theta_T) = \theta_0$$

where θ_0 and θ_T denote the values of θ at times zero and T , respectively. To understand this result, note that over a very small time interval the change in θ is normally distributed with zero mean. The expected change in θ over any very small time interval is therefore zero. The change in θ between time 0 and time T is the sum of its changes over many small time intervals. It follows that the expected change in θ between time 0 and time T must also be zero.

The Equivalent Martingale Measure Result

Suppose that f and g are the prices of traded securities dependent on a single source of uncertainty. Assume that the securities provide no income during the time period under consideration and define $\phi = f/g$.⁴ The variable ϕ is the relative price of f with respect to g . It can be thought of as measuring the price of f in units of g rather than dollars. The security price g is referred to as the *numeraire*.

² See S.A. Ross, "The Arbitrage Theory of Capital Asset Pricing," *Journal of Economic Theory*, 13 (December 1976): 343–62.

³ More formally, a sequence of random variables X_0, X_1, \dots is a martingale if $E(X_i | X_{i-1}, X_{i-2}, \dots, X_0) = X_{i-1}$, for all $i > 0$, where E denotes expectation.

⁴ Problem 27.8 extends the analysis to situations where the securities provide income.

The *equivalent martingale measure* result shows that, when there are no arbitrage opportunities, ϕ is a martingale for some choice of the market price of risk. What is more, for a given numeraire security g , the same choice of the market price of risk makes ϕ a martingale for all securities f . This choice of the market price of risk is the volatility of g . In other words, when the market price of risk is set equal to the volatility of g , the ratio f/g is a martingale for all security prices f . (Note that the market price of risk has the same dimension as volatility. Both are “per square root of time.” The choice for the market price of risk is therefore valid.)

To prove this result, suppose that the volatilities of f and g are σ_f and σ_g . From equation (27.10), in a world where the market price of risk is σ_g ,

$$df = (r + \sigma_g \sigma_f) f dt + \sigma_f f dz$$

$$dg = (r + \sigma_g^2) g dt + \sigma_g g dz$$

Using Itô's lemma gives

$$d \ln f = (r + \sigma_g \sigma_f - \sigma_f^2/2) dt + \sigma_f dz$$

$$d \ln g = (r + \sigma_g^2/2) dt + \sigma_g dz$$

so that

$$d(\ln f - \ln g) = (\sigma_g \sigma_f - \sigma_f^2/2 - \sigma_g^2/2) dt + (\sigma_f - \sigma_g) dz$$

or

$$d\left(\ln \frac{f}{g}\right) = -\frac{(\sigma_f - \sigma_g)^2}{2} dt + (\sigma_f - \sigma_g) dz$$

Itô's lemma can be used to determine the process for f/g from the process for $\ln(f/g)$:

$$d\left(\frac{f}{g}\right) = (\sigma_f - \sigma_g) \frac{f}{g} dz \quad (27.14)$$

This shows that f/g is a martingale and proves the equivalent martingale measure result. We will refer to a world where the market price of risk is the volatility σ_g of g as a world that is *forward risk neutral* with respect to g .

Because f/g is a martingale in a world that is forward risk neutral with respect to g , it follows from the result at the beginning of this section that

$$\frac{f_0}{g_0} = E_g\left(\frac{f_T}{g_T}\right)$$

or

$$f_0 = g_0 E_g\left(\frac{f_T}{g_T}\right) \quad (27.15)$$

where E_g denotes the expected value in a world that is forward risk neutral with respect to g .

27.4 ALTERNATIVE CHOICES FOR THE NUMERAIRE

We now present a number of examples of the equivalent martingale measure result. The first example shows that it is consistent with the traditional risk-neutral valuation result

used in earlier chapters. The other examples prepare the way for the valuation of bond options, interest rate caps, and swap options in Chapter 28.

Money Market Account as the Numeraire

The dollar money market account is a security that is worth \$1 at time zero and earns the instantaneous risk-free rate r at any given time.⁵ The variable r may be stochastic. If we set g equal to the money market account, it grows at rate r so that

$$dg = rg dt \quad (27.16)$$

The drift of g is stochastic, but the volatility of g is zero. It follows from the results in Section 27.3 that f/g is a martingale in a world where the market price of risk is zero. This is the world we defined earlier as the traditional risk-neutral world. From equation (27.15),

$$f_0 = g_0 \hat{E} \left(\frac{f_T}{g_T} \right) \quad (27.17)$$

where \hat{E} denotes expectations in the traditional risk-neutral world.

In this case, $g_0 = 1$ and

$$g_T = e^{\int_0^T r dt}$$

so that equation (27.17) reduces to

$$f_0 = \hat{E} \left(e^{-\int_0^T r dt} f_T \right) \quad (27.18)$$

or

$$f_0 = \hat{E} \left(e^{-\bar{r}T} f_T \right) \quad (27.19)$$

where \bar{r} is the average value of r between time 0 and time T . This equation shows that one way of valuing an interest rate derivative is to simulate the short-term interest rate r in the traditional risk-neutral world. On each trial the expected payoff is calculated and discounted at the average value of the short rate on the sampled path.

When the short-term interest rate r is assumed to be constant, equation (27.19) reduces to

$$f_0 = e^{-rT} \hat{E}(f_T)$$

or the risk-neutral valuation relationship used in earlier chapters.

Zero-Coupon Bond Price as the Numeraire

Define $P(t, T)$ as the price at time t of a zero-coupon bond that pays off \$1 at time T . We now explore the implications of setting g equal to $P(t, T)$. Let E_T denote expectations in a world that is forward risk neutral with respect to $P(t, T)$. Because $g_T = P(T, T) = 1$ and $g_0 = P(0, T)$, equation (27.15) gives

$$f_0 = P(0, T) E_T(f_T) \quad (27.20)$$

⁵ The money account is the limit as Δt approaches zero of the following security. For the first short period of time of length Δt , it is invested at the initial Δt period rate; at time Δt , it is reinvested for a further period of time Δt at the new Δt period rate; at time $2\Delta t$, it is again reinvested for a further period of time Δt at the new Δt period rate; and so on. The money market accounts in other currencies are defined analogously to the dollar money market account.

Notice the difference between equations (27.20) and (27.19). In equation (27.19), the discounting is inside the expectations operator. In equation (27.20) the discounting, as represented by the $P(0, T)$ term, is outside the expectations operator. The use of $P(t, T)$ as the numeraire therefore considerably simplifies things for a security that provides a payoff solely at time T .

Consider any variable θ that is not an interest rate.⁶ A forward contract on θ with maturity T is defined as a contract that pays off $\theta_T - K$ at time T , where θ_T is the value θ at time T . Define f as the value of this forward contract. From equation (27.20),

$$f_0 = P(0, T)[E_T(\theta_T) - K]$$

The forward price, F , of θ is the value of K for which f_0 equals zero. It therefore follows that

$$P(0, T)[E_T(\theta_T) - F] = 0$$

or

$$F = E_T(\theta_T) \quad (27.21)$$

Equation (27.21) shows that the forward price of any variable (except an interest rate) is its expected future spot price in a world that is forward risk neutral with respect to $P(t, T)$. Note the difference here between forward prices and futures prices. The argument in Section 17.7 shows that the futures price of a variable is the expected future spot price in the traditional risk-neutral world.

Equation (27.20) shows that any security that provides a payoff at time T can be valued by calculating its expected payoff in a world that is forward risk neutral with respect to a bond maturing at time T and discounting at the risk-free rate for maturity T . Equation (27.21) shows that it is correct to assume that the expected value of the underlying variables equal their forward values when computing the expected payoff.

Interest Rates When Zero-Coupon Bond Price is the Numeraire

For the next result, define $R(t, T, T^*)$ as the forward interest rate as seen at time t for the period between T and T^* expressed with a compounding period of $T^* - T$. (For example, if $T^* - T = 0.5$, the interest rate is expressed with semiannual compounding; if $T^* - T = 0.25$, it is expressed with quarterly compounding; and so on.) The forward price, as seen at time t , of a zero-coupon bond lasting between times T and T^* is

$$\frac{P(t, T^*)}{P(t, T)}$$

A forward interest rate is defined differently from the forward value of most variables. A forward interest rate is the interest rate implied by the corresponding forward bond price. It follows that

$$\frac{1}{[1 + (T^* - T)R(t, T, T^*)]} = \frac{P(t, T^*)}{P(t, T)}$$

so that

$$R(t, T, T^*) = \frac{1}{T^* - T} \left[\frac{P(t, T)}{P(t, T^*)} - 1 \right]$$

⁶ The analysis given here does not apply to interest rates because forward contracts for interest rates are defined differently from forward contracts for other variables. A forward interest rate is the interest rate implied by the corresponding forward bond price.

or

$$R(t, T, T^*) = \frac{1}{T^* - T} \left[\frac{P(t, T) - P(t, T^*)}{P(t, T^*)} \right]$$

Setting

$$f = \frac{1}{T^* - T} [P(t, T) - P(t, T^*)]$$

and $g = P(t, T^*)$, the equivalent martingale measure result shows that $R(t, T, T^*)$ is a martingale in a world that is forward risk neutral with respect to $P(t, T^*)$. This means that

$$R(0, T, T^*) = E_{T^*}[R(T, T, T^*)] \quad (27.22)$$

where E_{T^*} denotes expectations in a world that is forward risk neutral with respect to $P(t, T^*)$.

The variable $R(0, T, T^*)$ is the forward interest rate between times T and T^* as seen at time 0, whereas $R(T, T, T^*)$ is the realized interest rate between times T and T^* . Equation (27.22) therefore shows that the forward interest rate between times T and T^* equals the expected future interest rate in a world that is forward risk neutral with respect to a zero-coupon bond maturing at time T^* . This result, when combined with that in equation (27.20), will be critical to an understanding of the standard market model for interest rate caps in the next chapter.

Annuity Factor as the Numeraire

For the next application of equivalent martingale measure arguments, consider a swap starting at a future time T with payment dates at times T_1, T_2, \dots, T_N . Define $T_0 = T$. Assume that the principal underlying the swap is \$1. Suppose that the forward swap rate (i.e., the interest rate on the fixed side that makes the swap have a value of zero) is $s(t)$ at time t ($t \leq T$). The value of the fixed side of the swap is

$$s(t)A(t)$$

where

$$A(t) = \sum_{i=0}^{N-1} (T_{i+1} - T_i) P(t, T_{i+1})$$

Chapter 7 showed that, when the principal is added to the payment on the last payment date of a swap, the value of the floating side of the swap on the initiation date equals the underlying principal. It follows that if \$1 is added at time T_N , the floating side is worth \$1 at time T_0 . (This is because, when the discount rate is the LIBOR/swap rate, the present value of the payments on a LIBOR floating-rate bond equals the bond's principal.) The value of \$1 received at time T_N is $P(t, T_N)$. The value of \$1 at time T_0 is $P(t, T_0)$. The value of the floating side at time t is, therefore,

$$P(t, T_0) - P(t, T_N)$$

Equating the values of the fixed and floating sides gives

$$s(t)A(t) = P(t, T_0) - P(t, T_N)$$

or

$$s(t) = \frac{P(t, T_0) - P(t, T_N)}{A(t)} \quad (27.23)$$

The equivalent martingale measure result can be applied by setting f equal to $P(t, T_0) - P(t, T_N)$ and g equal to $A(t)$. This leads to

$$s(t) = E_A[s(T)] \quad (27.24)$$

where E_A denotes expectations in a world that is forward risk neutral with respect to $A(t)$. Therefore, in a world that is forward risk neutral with respect to $A(t)$, the expected future swap rate is the current swap rate.

For any security, f , the result in equation (27.15) shows that

$$f_0 = A(0)E_A\left[\frac{f_T}{A(T)}\right] \quad (27.25)$$

This result, when combined with the result in equation (27.24), will be critical to an understanding of the standard market model for European swap options in the next chapter.

27.5 EXTENSION TO SEVERAL FACTORS

The results presented in Sections 27.3 and 27.4 can be extended to cover the situation when there are many independent factors.⁷ Assume that there are n independent factors and that the processes for f and g in the traditional risk-neutral world are

$$df = rf dt + \sum_{i=1}^n \sigma_{f,i} f dz_i$$

and

$$dg = rg dt + \sum_{i=1}^n \sigma_{g,i} g dz_i$$

It follows from Section 27.2 that other internally consistent worlds can be defined by setting

$$df = \left[r + \sum_{i=1}^n \lambda_i \sigma_{f,i} \right] f dt + \sum_{i=1}^n \sigma_{f,i} f dz_i$$

and

$$dg = \left[r + \sum_{i=1}^n \lambda_i \sigma_{g,i} \right] g dt + \sum_{i=1}^n \sigma_{g,i} g dz_i$$

where the λ_i ($1 \leq i \leq n$) are the n market prices of risk. One of these other worlds is the real world.

The definition of forward risk neutrality can be extended so that a world is forward risk neutral with respect to g , where $\lambda_i = \sigma_{g,i}$ for all i . It can be shown from Itô's lemma, using the fact that the dz_i are uncorrelated, that the process followed by f/g in

⁷ The independence condition is not critical. If factors are not independent they can be orthogonalized.

this world has zero drift (see Problem 27.12). The rest of the results in the last two sections (from equation (27.15) onward) are therefore still true.

27.6 BLACK'S MODEL REVISITED

Section 17.8 explained that Black's model is a popular tool for pricing European options in terms of the forward or futures price of the underlying asset when interest rates are constant. We are now in a position to relax the constant interest rate assumption and show that Black's model can be used to price European options in terms of the forward price of the underlying asset when interest rates are stochastic.

Consider a European call option on an asset with strike price K that lasts until time T . From equation (27.20), the option's price is given by

$$c = P(0, T)E_T[\max(S_T - K, 0)] \quad (27.26)$$

where S_T is the asset price at time T and E_T denotes expectations in a world that is forward risk neutral with respect to $P(t, T)$. Define F_0 and F_T as the forward price of the asset at time 0 and time T for a contract maturing at time T . Because $S_T = F_T$,

$$c = P(0, T)E_T[\max(F_T - K, 0)]$$

Assume that F_T is lognormal in the world being considered, with the standard deviation of $\ln(F_T)$ equal to $\sigma_F\sqrt{T}$. This could be because the forward price follows a stochastic process with constant volatility σ_F . The appendix at the end of Chapter 14 shows that

$$E_T[\max(F_T - K, 0)] = E_T(F_T)N(d_1) - KN(d_2) \quad (27.27)$$

where

$$d_1 = \frac{\ln[E_T(F_T)/K] + \sigma_F^2 T/2}{\sigma_F\sqrt{T}}$$

$$d_2 = \frac{\ln[E_T(F_T)/K] - \sigma_F^2 T/2}{\sigma_F\sqrt{T}}$$

From equation (27.21), $E_T(F_T) = E_T(S_T) = F_0$. Hence,

$$c = P(0, T)[F_0N(d_1) - KN(d_2)] \quad (27.28)$$

where

$$d_1 = \frac{\ln[F_0/K] + \sigma_F^2 T/2}{\sigma_F\sqrt{T}}$$

$$d_2 = \frac{\ln[F_0/K] - \sigma_F^2 T/2}{\sigma_F\sqrt{T}}$$

Similarly,

$$p = P(0, T)[KN(-d_2) - F_0N(-d_1)] \quad (27.29)$$

where p is the price of a European put option on the asset with strike price K and time to maturity T . This is Black's model. It applies to both investment and consumption assets and, as we have just shown, is true when interest rates are stochastic provided that F_0 is the forward asset price. The variable σ_F can be interpreted as the (constant) volatility of the forward asset price.

27.7 OPTION TO EXCHANGE ONE ASSET FOR ANOTHER

Consider next an option to exchange an investment asset worth U for an investment asset worth V . This has already been discussed in Section 25.13. Suppose that the volatilities of U and V are σ_U and σ_V and the coefficient of correlation between them is ρ .

Assume first that the assets provide no income and choose the numeraire security g to be U . Setting $f = V$ in equation (27.15) gives

$$V_0 = U_0 E_U \left(\frac{V_T}{U_T} \right) \quad (27.30)$$

where E_U denotes expectations in a world that is forward risk neutral with respect to U .

The variable f in equation (27.15) can be set equal to the value of the option under consideration, so that $f_T = \max(V_T - U_T, 0)$. It follows that

$$f_0 = U_0 E_U \left[\frac{\max(V_T - U_T, 0)}{U_T} \right]$$

or

$$f_0 = U_0 E_U \left[\max \left(\frac{V_T}{U_T} - 1, 0 \right) \right] \quad (27.31)$$

The volatility of V/U is $\hat{\sigma}$ (see Problem 27.14), where

$$\hat{\sigma}^2 = \sigma_U^2 + \sigma_V^2 - 2\rho\sigma_U\sigma_V$$

From the appendix at the end of Chapter 14, equation (27.31) becomes

$$f_0 = U_0 \left[E_U \left(\frac{V_T}{U_T} \right) N(d_1) - N(d_2) \right]$$

where

$$d_1 = \frac{\ln(V_0/U_0) + \hat{\sigma}^2 T/2}{\hat{\sigma}\sqrt{T}} \quad \text{and} \quad d_2 = d_1 - \hat{\sigma}\sqrt{T}$$

Substituting from equation (27.30) gives

$$f_0 = V_0 N(d_1) - U_0 N(d_2) \quad (27.32)$$

This is the value of an option to exchange one asset for another when the assets provide no income.

Problem 27.8 shows that, when f and g provide income at rate q_f and q_g , equation (27.15) becomes

$$f_0 = g_0 e^{(q_f - q_g)T} E_g \left(\frac{f_T}{g_T} \right)$$

This means that equations (27.30) and (27.31) become

$$E_U \left(\frac{V_T}{U_T} \right) = e^{(q_U - q_V)T} \frac{V_0}{U_0}$$

and

$$f_0 = e^{-q_U T} U_0 E_U \left[\max \left(\frac{V_T}{U_T} - 1, 0 \right) \right]$$

and equation (27.32) becomes

$$f_0 = e^{-q_V T} V_0 N(d_1) - e^{-q_U T} U_0 N(d_2)$$

with d_1 and d_2 being redefined as

$$d_1 = \frac{\ln(V_0/U_0) + (q_U - q_V + \hat{\sigma}^2/2)T}{\hat{\sigma}\sqrt{T}} \quad \text{and} \quad d_2 = d_1 - \hat{\sigma}\sqrt{T}$$

This is the result given in equation (25.5) for the value of an option to exchange one asset for another.

27.8 CHANGE OF NUMERAIRE

In this section, we consider the impact of a change in numeraire on the process followed by a market variable. Suppose first that the variable is the price of a traded security, f . In a world where the market price of dz_i risk is λ_i ,

$$df = \left[r + \sum_{i=1}^n \lambda_i \sigma_{f,i} \right] f dt + \sum_{i=1}^n \sigma_{f,i} f dz_i$$

Similarly, when it is λ_i^* ,

$$df = \left[r + \sum_{i=1}^n \lambda_i^* \sigma_{f,i} \right] f dt + \sum_{i=1}^n \sigma_{f,i} f dz_i$$

The effect of moving from the first world to the second is therefore to increase the expected growth rate of the price of any traded security f by

$$\sum_{i=1}^n (\lambda_i^* - \lambda_i) \sigma_{f,i}$$

Consider next a variable v that is not the price of a traded security. As shown in Technical Note 20 at www.rotman.utoronto.ca/~hull/TechnicalNotes, the expected growth rate of v responds to a change in the market price of risk in the same way as the expected growth rate of the prices of traded securities. It increases by

$$\alpha_v = \sum_{i=1}^n (\lambda_i^* - \lambda_i) \sigma_{v,i} \tag{27.33}$$

where $\sigma_{v,i}$ is the i th component of the volatility of v .

When we move from a numeraire of g to a numeraire of h , $\lambda_i = \sigma_{g,i}$ and $\lambda_i^* = \sigma_{h,i}$. Define $w = h/g$ and $\sigma_{w,i}$ as the i th component of the volatility of w . From Itô's

lemma (see Problem 27.14),

$$\sigma_{w,i} = \sigma_{h,i} - \sigma_{g,i}$$

so that equation (27.33) becomes

$$\alpha_v = \sum_{i=1}^n \sigma_{w,i} \sigma_{v,i} \quad (27.34)$$

We will refer to w as the *numeraire ratio*. Equation (27.34) is equivalent to

$$\alpha_v = \rho \sigma_v \sigma_w \quad (27.35)$$

where σ_v is the total volatility of v , σ_w is the total volatility of w , and ρ is the instantaneous correlation between changes in v and w .⁸

This is a surprisingly simple result. The adjustment to the expected growth rate of a variable v when we change from one numeraire to another is the instantaneous covariance between the percentage change in v and the percentage change in the numeraire ratio. This result will be used when timing and quanto adjustments are considered in Chapter 29.

A particular case of the results in this section is when we move from the real world to the traditional risk-neutral world (where all the market prices of risk are zero). From equation (27.33), the growth rate of v changes by $-\sum_{i=1}^n \lambda_i \sigma_{vi}$. This corresponds to the result in equation (27.13) when v is the price of a traded security. We have shown that it is also true when v is not the price of a traded security. In general, the way that we move from one world to another for variables that are not the prices of traded securities are the same as for those that are.

SUMMARY

The market price of risk of a variable defines the trade-offs between risk and return for traded securities dependent on the variable. When there is one underlying variable, a derivative's excess return over the risk-free rate equals the market price of risk multiplied by the derivative's volatility. When there are many underlying variables, the excess return is the sum of the market price of risk multiplied by the volatility for each variable.

A powerful tool in the valuation of derivatives is risk-neutral valuation. This was introduced in Chapters 12 and 14. The principle of risk-neutral valuation shows that, if we assume that the world is risk neutral when valuing derivatives, we get the right answer—not just in a risk-neutral world, but in all other worlds as well. In the traditional risk-neutral world, the market price of risk of all variables is zero. This chapter has extended the principle of risk-neutral valuation. It has shown that, when

⁸ To see this, note that the changes Δv and Δw in v and w in a short period of time Δt are given by

$$\begin{aligned} \Delta v &= \dots + \sum \sigma_{v,i} v \epsilon_i \sqrt{\Delta t} \\ \Delta w &= \dots + \sum \sigma_{w,i} w \epsilon_i \sqrt{\Delta t} \end{aligned}$$

Since the dz_i are uncorrelated, it follows that $E(\epsilon_i \epsilon_j) = 0$ when $i \neq j$. Also, from the definition of ρ , we have

$$\rho \sigma_v \sigma_w = E(\Delta v \Delta w) - E(\Delta v) E(\Delta w)$$

When terms of higher order than Δt are ignored this leads to

$$\rho \sigma_v \sigma_w = \sum \sigma_{w,i} \sigma_{v,i}$$

interest rates are stochastic, there are many interesting and useful alternatives to the traditional risk-neutral world.

A martingale is a zero drift stochastic process. Any variable following a martingale has the simplifying property that its expected value at any future time equals its value today. The equivalent martingale measure result shows that, if g is a security price, there is a world in which the ratio f/g is a martingale for all security prices f . It turns out that, by appropriately choosing the numeraire security g , the valuation of many interest rate dependent derivatives can be simplified.

This chapter has used the equivalent martingale measure result to extend Black's model to the situation where interest rates are stochastic and to value an option to exchange one asset for another. In Chapters 28 to 32, it will be useful in valuing interest rate derivatives.

FURTHER READING

Baxter, M., and A. Rennie, *Financial Calculus*. Cambridge University Press, 1996.

Cox, J. C., J. E. Ingersoll, and S. A. Ross, "An Intertemporal General Equilibrium Model of Asset Prices," *Econometrica*, 53 (1985): 363–84.

Duffie, D., *Dynamic Asset Pricing Theory*, 3rd edn. Princeton University Press, 2001.

Garman, M., "A General Theory of Asset Valuation Under Diffusion State Processes," Working Paper 50, University of California, Berkeley, 1976.

Harrison, J. M., and D. M. Kreps, "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory*, 20 (1979): 381–408.

Harrison, J. M., and S. R. Pliska, "Martingales and Stochastic Integrals in the Theory of Continuous Trading," *Stochastic Processes and Their Applications*, 11 (1981): 215–60.

Practice Questions (Answers in the Solutions Manual)

- 27.1. How is the market price of risk defined for a variable that is not the price of an investment asset?
- 27.2. Suppose that the market price of risk for gold is zero. If the storage costs are 1% per annum and the risk-free rate of interest is 6% per annum, what is the expected growth rate in the price of gold? Assume that gold provides no income.
- 27.3. Consider two securities both of which are dependent on the same market variable. The expected returns from the securities are 8% and 12%. The volatility of the first security is 15%. The instantaneous risk-free rate is 4%. What is the volatility of the second security?
- 27.4. An oil company is set up solely for the purpose of exploring for oil in a certain small area of Texas. Its value depends primarily on two stochastic variables: the price of oil and the quantity of proven oil reserves. Discuss whether the market price of risk for the second of these two variables is likely to be positive, negative, or zero.
- 27.5. Deduce the differential equation for a derivative dependent on the prices of two non-dividend-paying traded securities by forming a riskless portfolio consisting of the derivative and the two traded securities.

- 27.6. Suppose that an interest rate x follows the process

$$dx = a(x_0 - x) dt + c \sqrt{x} dz$$

where a , x_0 , and c are positive constants. Suppose further that the market price of risk for x is λ . What is the process for x in the traditional risk-neutral world?

- 27.7. Prove that, when the security f provides income at rate q , equation (27.9) becomes $\mu + q - r = \lambda\sigma$. (*Hint*: Form a new security f^* that provides no income by assuming that all the income from f is reinvested in f .)
- 27.8. Show that when f and g provide income at rates q_f and q_g , respectively, equation (27.15) becomes

$$f_0 = g_0 e^{(q_f - q_g)T} E_g \left(\frac{f_T}{g_T} \right)$$

(*Hint*: Form new securities f^* and g^* that provide no income by assuming that all the income from f is reinvested in f and all the income in g is reinvested in g .)

- 27.9. “The expected future value of an interest rate in a risk-neutral world is greater than it is in the real world.” What does this statement imply about the market price of risk for (a) an interest rate and (b) a bond price. Do you think the statement is likely to be true? Give reasons.
- 27.10. The variable S is an investment asset providing income at rate q measured in currency A. It follows the process

$$dS = \mu_S S dt + \sigma_S S dz$$

in the real world. Defining new variables as necessary, give the process followed by S , and the corresponding market price of risk, in:

- A world that is the traditional risk-neutral world for currency A
 - A world that is the traditional risk-neutral world for currency B
 - A world that is forward risk neutral with respect to a zero-coupon currency A bond maturing at time T
 - A world that is forward risk neutral with respect to a zero coupon currency B bond maturing at time T .
- 27.11. Explain the difference between the way a forward interest rate is defined and the way the forward values of other variables such as stock prices, commodity prices, and exchange rates are defined.
- 27.12. Prove the result in Section 27.5 that when

$$df = \left[r + \sum_{i=1}^n \lambda_i \sigma_{f,i} \right] f dt + \sum_{i=1}^n \sigma_{f,i} f dz_i$$

and

$$dg = \left[r + \sum_{i=1}^n \lambda_i \sigma_{g,i} \right] g dt + \sum_{i=1}^n \sigma_{g,i} g dz_i$$

with the dz_i uncorrelated, f/g is a martingale for $\lambda_i = \sigma_{g,i}$. (*Hint*: Start by using equation (13A.11) to get the processes for $\ln f$ and $\ln g$.)

- 27.13. Show that when $w = h/g$ and h and g are each dependent on n Wiener processes, the i th component of the volatility of w is the i th component of the volatility of h minus the i th component of the volatility of g . (*Hint*: Start by using equation (13A.11) to get the processes for $\ln f$ and $\ln g$.)

Further Questions

27.14. A security's price is positively dependent on two variables: the price of copper and the yen/dollar exchange rate. Suppose that the market price of risk for these variables is 0.5 and 0.1, respectively. If the price of copper were held fixed, the volatility of the security would be 8% per annum; if the yen/dollar exchange rate were held fixed, the volatility of the security would be 12% per annum. The risk-free interest rate is 7% per annum. What is the expected rate of return from the security? If the two variables are uncorrelated with each other, what is the volatility of the security?

27.15. Suppose that the price of a zero-coupon bond maturing at time T follows the process

$$dP(t, T) = \mu_P P(t, T) dt + \sigma_P P(t, T) dz$$

and the price of a derivative dependent on the bond follows the process

$$df = \mu_f f dt + \sigma_f f dz$$

Assume only one source of uncertainty and that f provides no income.

- What is the forward price F of f for a contract maturing at time T ?
- What is the process followed by F in a world that is forward risk neutral with respect to $P(t, T)$?
- What is the process followed by F in the traditional risk-neutral world?
- What is the process followed by f in a world that is forward risk neutral with respect to a bond maturing at time T^* , where $T^* \neq T$? Assume that σ_P^* is the volatility of this bond.

27.16. Consider a variable that is not an interest rate:

- In what world is the futures price of the variable a martingale?
- In what world is the forward price of the variable a martingale?
- Defining variables as necessary, derive an expression for the difference between the drift of the futures price and the drift of the forward price in the traditional risk-neutral world.
- Show that your result is consistent with the points made in Section 5.8 about the circumstances when the futures price is above the forward price.



CHAPTER 28

Interest Rate Derivatives: The Standard Market Models

Interest rate derivatives are instruments whose payoffs are dependent in some way on the level of interest rates. In the 1980s and 1990s, the volume of trading in interest rate derivatives in both the over-the-counter and exchange-traded markets increased rapidly. Many new products were developed to meet particular needs of end users. A key challenge for derivatives traders was to find good, robust procedures for pricing and hedging these products. Interest rate derivatives are more difficult to value than equity and foreign exchange derivatives for the following reasons:

1. The behavior of an individual interest rate is more complicated than that of a stock price or an exchange rate.
2. For the valuation of many products it is necessary to develop a model describing the behavior of the entire zero-coupon yield curve.
3. The volatilities of different points on the yield curve are different.
4. Interest rates are used for discounting the derivative as well as defining its payoff.

This chapter considers the three most popular over-the-counter interest rate option products: bond options, interest rate caps/floors, and swap options. It explains how the products work and the standard market models used to value them.

28.1 BOND OPTIONS

A bond option is an option to buy or sell a particular bond by a particular date for a particular price. In addition to trading in the over-the-counter market, bond options are frequently embedded in bonds when they are issued to make them more attractive to either the issuer or potential purchasers.

Embedded Bond Options

One example of a bond with an embedded bond option is a *callable bond*. This is a bond that contains provisions allowing the issuing firm to buy back the bond at a

predetermined price at certain times in the future. The holder of such a bond has sold a call option to the issuer. The strike price or call price in the option is the predetermined price that must be paid by the issuer to the holder. Callable bonds cannot usually be called for the first few years of their life. (This is known as the lock-out period.) After that, the call price is usually a decreasing function of time. For example, in a 10-year callable bond, there might be no call privileges for the first 2 years. After that, the issuer might have the right to buy the bond back at a price of 110 in years 3 and 4 of its life, at a price of 107.5 in years 5 and 6, at a price of 106 in years 7 and 8, and at a price of 103 in years 9 and 10. The value of the call option is reflected in the quoted yields on bonds. Bonds with call features generally offer higher yields than bonds with no call features.

Another type of bond with an embedded option is a *puttable bond*. This contains provisions that allow the holder to demand early redemption at a predetermined price at certain times in the future. The holder of such a bond has purchased a put option on the bond as well as the bond itself. Because the put option increases the value of the bond to the holder, bonds with put features provide lower yields than bonds with no put features. A simple example of a puttable bond is a 10-year bond where the holder has the right to be repaid at the end of 5 years. (This is sometimes referred to as a *retractable bond*.)

Loan and deposit instruments also often contain embedded bond options. For example, a 5-year fixed-rate deposit with a financial institution that can be redeemed without penalty at any time contains an American put option on a bond. (The deposit instrument is a bond that the investor has the right to put back to the financial institution at its face value at any time.) Prepayment privileges on loans and mortgages are similarly call options on bonds.

Finally, a loan commitment made by a bank or other financial institution is a put option on a bond. Consider, for example, the situation where a bank quotes a 5-year interest rate of 5% per annum to a potential borrower and states that the rate is good for the next 2 months. The client has, in effect, obtained the right to sell a 5-year bond with a 5% coupon to the financial institution for its face value any time within the next 2 months. The option will be exercised if rates increase.

European Bond Options

Many over-the-counter bond options and some embedded bond options are European. The assumption made in the standard market model for valuing European bond options is that the forward bond price has a constant volatility σ_B . This allows Black's model in Section 27.6 to be used. In equations (27.28) and (27.29), σ_F is set equal to σ_B and F_0 is set equal to the forward bond price F_B , so that

$$c = P(0, T)[F_B N(d_1) - KN(d_2)] \quad (28.1)$$

$$p = P(0, T)[KN(-d_2) - F_B N(-d_1)] \quad (28.2)$$

where

$$d_1 = \frac{\ln(F_B/K) + \sigma_B^2 T/2}{\sigma_B \sqrt{T}} \quad \text{and} \quad d_2 = d_1 - \sigma_B \sqrt{T}$$

with K the strike price of the bond option and T its time to maturity.

From Section 5.5, F_B can be calculated using the formula

$$F_B = \frac{B_0 - I}{P(0, T)} \quad (28.3)$$

where B_0 is the bond price at time zero and I is the present value of the coupons that will be paid during the life of the option. In this formula, both the spot bond price and the forward bond price are cash prices rather than quoted prices. The relationship between cash and quoted bond prices is explained in Section 6.1.

The strike price K in equations (28.1) and (28.2) should be the cash strike price. In choosing the correct value for K , the precise terms of the option are therefore important. If the strike price is defined as the cash amount that is exchanged for the bond when the option is exercised, K should be set equal to this strike price. If, as is more common, the strike price is the quoted price applicable when the option is exercised, K should be set equal to the strike price plus accrued interest at the expiration date of the option. Traders refer to the quoted price of a bond as the *clean price* and the cash price as the *dirty price*.

Example 28.1

Consider a 10-month European call option on a 9.75-year bond with a face value of \$1,000. (When the option matures, the bond will have 8 years and 11 months remaining.) Suppose that the current cash bond price is \$960, the strike price is \$1,000, the 10-month risk-free interest rate is 10% per annum, and the volatility of the forward bond price for a contract maturing in 10 months is 9% per annum. The bond pays a coupon of 10% per year (with payments made semiannually). Coupon payments of \$50 are expected in 3 months and 9 months. (This means that the accrued interest is \$25 and the quoted bond price is \$935.) We suppose that the 3-month and 9-month risk-free interest rates are 9.0% and 9.5% per annum, respectively. The present value of the coupon payments is, therefore,

$$50e^{-0.25 \times 0.09} + 50e^{-0.75 \times 0.095} = 95.45$$

or \$95.45. The bond forward price is from equation (28.3) given by

$$F_B = (960 - 95.45)e^{0.1 \times 0.8333} = 939.68$$

(a) If the strike price is the cash price that would be paid for the bond on exercise, the parameters for equation (28.1) are $F_B = 939.68$, $K = 1000$, $P(0, T) = e^{-0.1 \times (10/12)} = 0.9200$, $\sigma_B = 0.09$, and $T = 10/12$. The price of the call option is \$9.49.

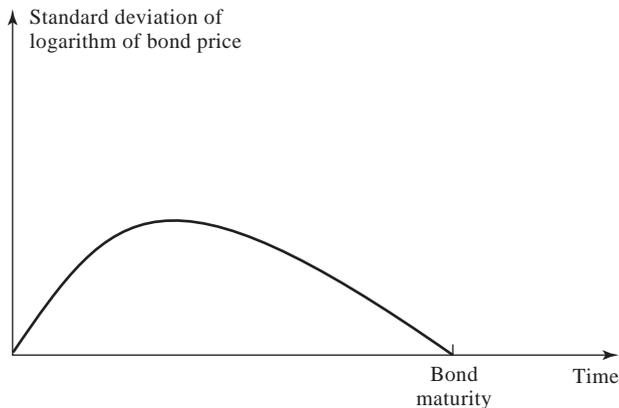
(b) If the strike price is the quoted price that would be paid for the bond on exercise, 1 month's accrued interest must be added to K because the maturity of the option is 1 month after a coupon date. This produces a value for K of

$$1,000 + 100 \times 0.08333 = 1,008.33$$

The values for the other parameters in equation (28.1) are unchanged (i.e., $F_B = 939.68$, $P(0, T) = 0.9200$, $\sigma_B = 0.09$, and $T = 0.8333$). The price of the option is \$7.97.

Figure 28.1 shows how the standard deviation of the logarithm of a bond's price

Figure 28.1 Standard deviation of logarithm of bond price at future times.



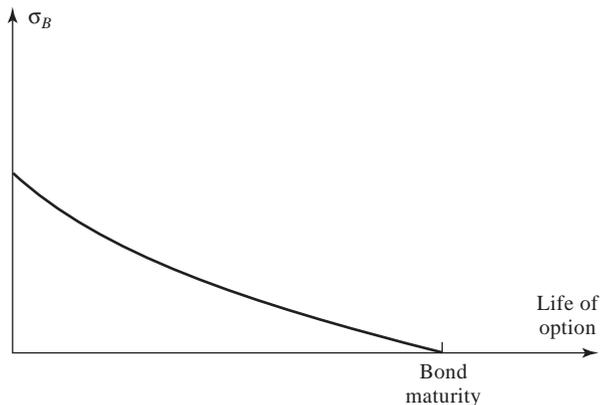
changes as we look further ahead. The standard deviation is zero today because there is no uncertainty about the bond’s price today. It is also zero at the bond’s maturity because we know that the bond’s price will equal its face value at maturity. Between today and the maturity of the bond, the standard deviation first increases and then decreases.

The volatility σ_B that should be used when a European option on the bond is valued is

$$\frac{\text{Standard deviation of logarithm of bond price at maturity of option}}{\sqrt{\text{Time to maturity of option}}}$$

What happens when, for a particular underlying bond, the life of the option is increased? Figure 28.2 shows a typical pattern for σ_B as a function of the life of the option, with σ_B declining as the life of the option increases.

Figure 28.2 Variation of forward bond price volatility σ_B with life of option when bond is kept fixed.



Yield Volatilities

The volatilities that are quoted for bond options are often yield volatilities rather than price volatilities. The duration concept, introduced in Chapter 4, is used by the market to convert a quoted yield volatility into a price volatility. Suppose that D is the modified duration of the bond underlying the option at the option maturity, as defined in Chapter 4. The relationship between the change ΔF_B in the forward bond price F_B and the change Δy_F in the forward yield y_F is

$$\frac{\Delta F_B}{F_B} \approx -D\Delta y_F$$

or

$$\frac{\Delta F_B}{F_B} \approx -Dy_F \frac{\Delta y_F}{y_F}$$

Volatility is a measure of the standard deviation of percentage changes in the value of a variable. This equation therefore suggests that the volatility of the forward bond price σ_B used in Black's model can be approximately related to the volatility of the forward bond yield σ_y by

$$\sigma_B = Dy_0\sigma_y \quad (28.4)$$

where y_0 is the initial value of y_F . When a yield volatility is quoted for a European bond option, the implicit assumption is usually that it will be converted to a price volatility using equation (28.4), and that this volatility will then be used in conjunction with equation (28.1) or (28.2) to obtain the option's price. Suppose that the bond underlying a call option will have a modified duration of 5 years at option maturity, the forward yield is 8%, and the forward yield volatility quoted by a broker is 20%. This means that the market price of the option corresponding to the broker quote is the price given by equation (28.1) when the volatility variable σ_B is

$$5 \times 0.08 \times 0.2 = 0.08$$

or 8% per annum. Figure 28.2 shows that forward bond volatilities depend on the option considered. Forward yield volatilities as we have just defined them are more constant. This is why traders prefer them.

The Bond_Options worksheet of the software DerivaGem accompanying this book can be used to price European bond options using Black's model by selecting Black-European as the Pricing Model. The user inputs a yield volatility, which is handled in the way just described. The strike price can be the cash or quoted strike price.

Example 28.2

Consider a European put option on a 10-year bond with a principal of 100. The coupon is 8% per year payable semiannually. The life of the option is 2.25 years and the strike price of the option is 115. The forward yield volatility is 20%. The zero curve is flat at 5% with continuous compounding. The DerivaGem software accompanying this book shows that the quoted price of the bond is 122.82. The price of the option when the strike price is a quoted price is \$2.36. When the strike price is a cash price, the price of the option is \$1.74. (See Problem 28.16 for the manual calculation.)

28.2 INTEREST RATE CAPS AND FLOORS

A popular interest rate option offered by financial institutions in the over-the-counter market is an *interest rate cap*. Interest rate caps can best be understood by first considering a floating-rate note where the interest rate is reset periodically equal to LIBOR. The time between resets is known as the *tenor*. Suppose the tenor is 3 months. The interest rate on the note for the first 3 months is the initial 3-month LIBOR rate; the interest rate for the next 3 months is set equal to the 3-month LIBOR rate prevailing in the market at the 3-month point; and so on.

An interest rate cap is designed to provide insurance against the rate of interest on the floating-rate note rising above a certain level. This level is known as the *cap rate*. Suppose that the principal amount is \$10 million, the tenor is 3 months, the life of the cap is 5 years, and the cap rate is 4%. (Because the payments are made quarterly, this cap rate is expressed with quarterly compounding.) The cap provides insurance against the interest on the floating rate note rising above 4%.

For the moment we ignore day count issues and assume that there is exactly 0.25 year between each payment date. (We will cover day count issues at the end of this section.) Suppose that on a particular reset date the 3-month LIBOR interest rate is 5%. The floating rate note would require

$$0.25 \times 0.05 \times \$10,000,000 = \$125,000$$

of interest to be paid 3 months later. With a 3-month LIBOR rate of 4% the interest payment would be

$$0.25 \times 0.04 \times \$10,000,000 = \$100,000$$

The cap therefore provides a payoff of \$25,000. The payoff does not occur on the reset date when the 5% is observed: it occurs 3 months later. This reflects the usual time lag between an interest rate being observed and the corresponding payment being required.

At each reset date during the life of the cap, LIBOR is observed. If LIBOR is less than 4%, there is no payoff from the cap three months later. If LIBOR is greater than 4%, the payoff is one quarter of the excess applied to the principal of \$10 million. Note that caps are usually defined so that the initial LIBOR rate, even if it is greater than the cap rate, does not lead to a payoff on the first reset date. In our example, the cap lasts for 5 years. There are, therefore, a total of 19 reset dates (at times 0.25, 0.50, 0.75, ..., 4.75 years) and 19 potential payoffs from the caps (at times 0.50, 0.75, 1.00, ..., 5.00 years).

The Cap as a Portfolio of Interest Rate Options

Consider a cap with a total life of T , a principal of L , and a cap rate of R_K . Suppose that the reset dates are t_1, t_2, \dots, t_n and define $t_{n+1} = T$. Define R_k as the LIBOR interest rate for the period between time t_k and t_{k+1} observed at time t_k ($1 \leq k \leq n$). The cap leads to a payoff at time t_{k+1} ($k = 1, 2, \dots, n$) of

$$L\delta_k \max(R_k - R_K, 0) \tag{28.5}$$

where $\delta_k = t_{k+1} - t_k$.¹ Both R_k and R_K are expressed with a compounding frequency equal to the frequency of resets.

¹ Day count issues are discussed at the end of this section.

Expression (28.5) is the payoff from a call option on the LIBOR rate observed at time t_k with the payoff occurring at time t_{k+1} . The cap is a portfolio of n such options. LIBOR rates are observed at times $t_1, t_2, t_3, \dots, t_n$ and the corresponding payoffs occur at times $t_2, t_3, t_4, \dots, t_{n+1}$. The n call options underlying the cap are known as *caplets*.

A Cap as a Portfolio of Bond Options

An interest rate cap can also be characterized as a portfolio of put options on zero-coupon bonds with payoffs on the puts occurring at the time they are calculated. The payoff in expression (28.5) at time t_{k+1} is equivalent to

$$\frac{L\delta_k}{1 + R_k\delta_k} \max(R_k - R_K, 0)$$

at time t_k . A few lines of algebra show that this reduces to

$$\max\left[L - \frac{L(1 + R_K\delta_k)}{1 + R_k\delta_k}, 0\right] \quad (28.6)$$

The expression

$$\frac{L(1 + R_K\delta_k)}{1 + R_k\delta_k}$$

is the value at time t_k of a zero-coupon bond that pays off $L(1 + R_K\delta_k)$ at time t_{k+1} . The expression in (28.6) is therefore the payoff from a put option with maturity t_k on a zero-coupon bond with maturity t_{k+1} when the face value of the bond is $L(1 + R_K\delta_k)$ and the strike price is L . It follows that an interest rate cap can be regarded as a portfolio of European put options on zero-coupon bonds.

Floors and Collars

Interest rate floors and interest rate collars (sometimes called floor–ceiling agreements) are defined analogously to caps. A *floor* provides a payoff when the interest rate on the underlying floating-rate note falls below a certain rate. With the notation already introduced, a floor provides a payoff at time t_{k+1} ($k = 1, 2, \dots, n$) of

$$L\delta_k \max(R_K - R_k, 0)$$

Analogously to an interest rate cap, an interest rate floor is a portfolio of put options on interest rates or a portfolio of call options on zero-coupon bonds. Each of the individual options comprising a floor is known as a *floorlet*. A *collar* is an instrument designed to guarantee that the interest rate on the underlying LIBOR floating-rate note always lies between two levels. A collar is a combination of a long position in a cap and a short position in a floor. It is usually constructed so that the price of the cap is initially equal to the price of the floor. The cost of entering into the collar is then zero.

Business Snapshot 28.1 gives the put–call parity relationship between caps and floors.

Valuation of Caps and Floors

As shown in equation (28.5), the caplet corresponding to the rate observed at time t_k provides a payoff at time t_{k+1} of

$$L\delta_k \max(R_k - R_K, 0)$$

Business Snapshot 28.1 Put–Call Parity for Caps and Floors

There is a put–call parity relationship between the prices of caps and floors. This is

$$\text{Value of cap} = \text{Value of floor} + \text{Value of swap}$$

In this relationship, the cap and floor have the same strike price, R_K . The swap is an agreement to receive LIBOR and pay a fixed rate of R_K with no exchange of payments on the first reset date. All three instruments have the same life and the same frequency of payments.

To see that the result is true, consider a long position in the cap combined with a short position in the floor. The cap provides a cash flow of $\text{LIBOR} - R_K$ for periods when LIBOR is greater than R_K . The short floor provides a cash flow of $-(R_K - \text{LIBOR}) = \text{LIBOR} - R_K$ for periods when LIBOR is less than R_K . There is therefore a cash flow of $\text{LIBOR} - R_K$ in all circumstances. This is the cash flow on the swap. It follows that the value of the cap minus the value of the floor must equal the value of the swap.

Note that swaps are usually structured so that LIBOR at time zero determines a payment on the first reset date. Caps and floors are usually structured so that there is no payoff on the first reset date. This is why put–call parity involves a nonstandard swap where there is no payment on the first reset date.

Under the standard market model, the value of the caplet is

$$L\delta_k P(0, t_{k+1})[F_k N(d_1) - R_K N(d_2)] \quad (28.7)$$

where

$$d_1 = \frac{\ln(F_k/R_K) + \sigma_k^2 t_k/2}{\sigma_k \sqrt{t_k}}$$

$$d_2 = \frac{\ln(F_k/R_K) - \sigma_k^2 t_k/2}{\sigma_k \sqrt{t_k}} = d_1 - \sigma_k \sqrt{t_k}$$

Here, F_k is the forward interest rate at time 0 for the period between time t_k and t_{k+1} , and σ_k is the volatility of this forward interest rate. This is a natural extension of Black's model. The volatility σ_k is multiplied by $\sqrt{t_k}$ because the interest rate R_k is observed at time t_k , but the discount factor $P(0, t_{k+1})$ reflects the fact that the payoff is at time t_{k+1} , not time t_k . The value of the corresponding floorlet is

$$L\delta_k P(0, t_{k+1})[R_K N(-d_2) - F_k N(-d_1)] \quad (28.8)$$

Example 28.3

Consider a contract that caps the LIBOR interest rate on \$10 million at 8% per annum (with quarterly compounding) for 3 months starting in 1 year. This is a caplet and could be one element of a cap. Suppose that the LIBOR/swap zero curve is flat at 7% per annum with quarterly compounding and the volatility of the 3-month forward rate underlying the caplet is 20% per annum. The continuously compounded zero rate for all maturities is 6.9395%. In equation (28.7), $F_k = 0.07$, $\delta_k = 0.25$, $L = 10$, $R_K = 0.08$, $t_k = 1.0$, $t_{k+1} = 1.25$, $P(0, t_{k+1}) =$

$e^{-0.069395 \times 1.25} = 0.9169$, and $\sigma_k = 0.20$. Also,

$$d_1 = \frac{\ln(0.07/0.08) + 0.2^2 \times 1/2}{0.20 \times 1} = -0.5677$$

$$d_2 = d_1 - 0.20 = -0.7677$$

so that the caplet price (in \$ millions) is

$$0.25 \times 10 \times 0.9169[0.07N(-0.5677) - 0.08N(-0.7677)] = \$0.005162$$

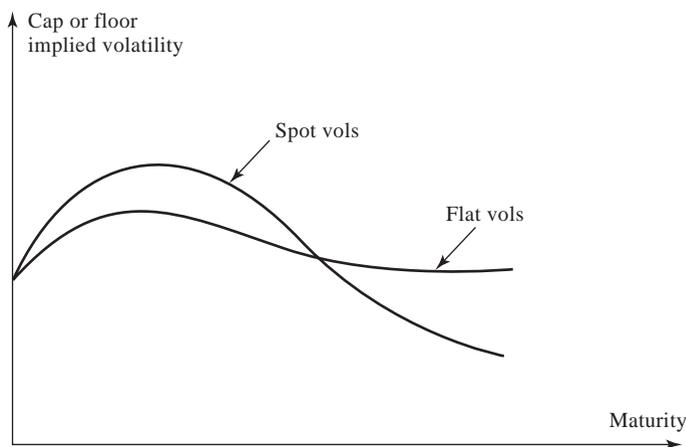
It is \$5,162. This result can also be obtained using the DerivaGem software accompanying this book.

Each caplet of a cap must be valued separately using equation (28.7). Similarly, each floorlet of a floor must be valued separately using equation (28.8). One approach is to use a different volatility for each caplet (or floorlet). The volatilities are then referred to as *spot volatilities*. An alternative approach is to use the same volatility for all the caplets (floorlets) comprising any particular cap (floor) but to vary this volatility according to the life of the cap (floor). The volatilities used are then referred to as *flat volatilities*.² The volatilities quoted in the market are usually flat volatilities. However, many traders like to estimate spot volatilities because this allows them to identify underpriced and overpriced caplets (floorlets). The put (call) options on Eurodollar futures are very similar to caplets (floorlets) and the spot volatilities used for caplets and floorlets on 3-month LIBOR are frequently compared with those calculated from the prices of Eurodollar futures options.

Spot Volatilities vs. Flat Volatilities

Figure 28.3 shows a typical pattern for spot volatilities and flat volatilities as a function of maturity. (In the case of a spot volatility, the maturity is the maturity of a caplet or floorlet; in the case of a flat volatility, it is the maturity of a cap or floor.) The flat

Figure 28.3 The volatility hump.



² Flat volatilities can be calculated from spot volatilities and vice versa (see Problem 28.20).

Table 28.1 Typical broker implied flat volatility quotes for US dollar caps and floors (% per annum).

<i>Life</i>	<i>Cap bid</i>	<i>Cap offer</i>	<i>Floor bid</i>	<i>Floor offer</i>
1 year	18.00	20.00	18.00	20.00
2 years	23.25	24.25	23.75	24.75
3 years	24.00	25.00	24.50	25.50
4 years	23.75	24.75	24.25	25.25
5 years	23.50	24.50	24.00	25.00
7 years	21.75	22.75	22.00	23.00
10 years	20.00	21.00	20.25	21.25

volatilities are akin to cumulative averages of the spot volatilities and therefore exhibit less variability. As indicated by Figure 28.3, a “hump” in the volatilities is usually observed. The peak of the hump is at about the 2- to 3-year point. This hump is observed both when the volatilities are implied from option prices and when they are calculated from historical data. There is no general agreement on the reason for the existence of the hump. One possible explanation is as follows. Rates at the short end of the zero curve are controlled by central banks. By contrast, 2- and 3-year interest rates are determined to a large extent by the activities of traders. These traders may be overreacting to the changes observed in the short rate and causing the volatility of these rates to be higher than the volatility of short rates. For maturities beyond 2 to 3 years, the mean reversion of interest rates, which is discussed in Chapter 30, causes volatilities to decline.

Interdealer brokers provide tables of implied flat volatilities for caps and floors. The instruments underlying the quotes are usually “at the money”. This is defined as the situation where the cap/floor rate equals the swap rate for a swap that has the same payment dates as the cap. Table 28.1 shows typical broker quotes for the US dollar market. The tenor of the cap is 3 months and the cap life varies from 1 to 10 years. The data exhibits the type of “hump” shown in Figure 28.3.

Theoretical Justification for the Model

The extension of Black’s model used to value a caplet can be shown to be internally consistent by considering a world that is forward risk neutral with respect to a zero-coupon bond maturing at time t_{k+1} . The analysis in Section 27.4 shows that:

1. The current value of any security is its expected value at time t_{k+1} in this world multiplied by the price of a zero-coupon bond maturing at time t_{k+1} (see equation (27.20)).
2. The expected value of an interest rate lasting between times t_k and t_{k+1} equals the forward interest rate in this world (see equation (27.22)).

The first of these results shows that, with the notation introduced earlier, the price of a caplet that provides a payoff at time t_{k+1} is

$$L\delta_k P(0, t_{k+1})E_{k+1}[\max(R_k - R_K, 0)] \quad (28.9)$$

where E_{k+1} denotes expected value in a world that is forward risk neutral with respect to a zero-coupon bond maturing at time t_{k+1} . When the forward interest rate underlying the cap (initially F_k) is assumed to have a constant volatility σ_k , R_k is lognormal in the world we are considering, with the standard deviation of $\ln(R_k)$ equal to $\sigma_k\sqrt{t_k}$. From the appendix at the end of Chapter 14, equation (28.9) becomes

$$L\delta_k P(0, t_{k+1})[E_{k+1}(R_k)N(d_1) - R_K N(d_2)]$$

where

$$d_1 = \frac{\ln[E_{k+1}(R_k)/R_K] + \sigma_k^2 t_k/2}{\sigma_k\sqrt{t_k}}$$

$$d_2 = \frac{\ln[E_{k+1}(R_k)/R_K] - \sigma_k^2 t_k/2}{\sigma_k\sqrt{t_k}} = d_1 - \sigma_k\sqrt{t_k}$$

The second result implies that

$$E_{k+1}(R_k) = F_k$$

Together the results lead to the cap pricing model in equation (28.7). They show that we can discount at the t_{k+1} -maturity interest rate observed in the market today providing we set the expected interest rate equal to the forward interest rate.

Use of DerivaGem

The software DerivaGem accompanying this book can be used to price interest rate caps and floors using Black's model. In the Cap_and_Swap_Option worksheet select Cap/Floor as the Underlying Type and Black-European as the Pricing Model. The zero curve is input using continuously compounded rates. The inputs include the start and end date of the period covered by the cap, the flat volatility, and the cap settlement frequency (i.e., the tenor). The software calculates the payment dates by working back from the end of period covered by the cap to the beginning. The initial caplet/floorlet is assumed to cover a period of length between 0.5 and 1.5 times a regular period. Suppose, for example, that the period covered by the cap is 1.22 years to 2.80 years and the settlement frequency is quarterly. There are six caplets covering the periods 2.55 to 2.80 years, 2.30 to 2.55 years, 2.05 to 2.30 years, 1.80 to 2.05 years, 1.55 to 1.80 years, and 1.22 to 1.55 years.

The Impact of Day Count Conventions

The formulas we have presented so far in this section do not reflect day count conventions (see Section 6.1 for an explanation of day count conventions). Suppose that the cap rate R_K is expressed with an actual/360 day count (as would be normal in the United States). This means that the time interval δ_k in the formulas should be replaced by a_k , the *accrual fraction* for the time period between t_k and t_{k+1} . Suppose, for example, that t_k is May 1 and t_{k+1} is August 1. Under actual/360 there are 92 days between these payment dates so that $a_k = 92/360 = 0.2556$. The forward rate F_k must be expressed with an actual/360 day count. This means that we must set it by solving

$$1 + a_k F_k = \frac{P(0, t_k)}{P(0, t_{k+1})}$$

The impact of all this is much the same as calculating δ_k on an actual/actual basis converting R_K from actual/360 to actual/actual, and calculating F_k on an actual/actual basis by solving

$$1 + \delta_k F_k = \frac{P(0, t_k)}{P(0, t_{k+1})}$$

28.3 EUROPEAN SWAP OPTIONS

Swap options, or *swaptions*, are options on interest rate swaps and are another popular type of interest rate option. They give the holder the right to enter into a certain interest rate swap at a certain time in the future. (The holder does not, of course, have to exercise this right.) Many large financial institutions that offer interest rate swap contracts to their corporate clients are also prepared to sell them swaptions or buy swaptions from them. As shown in Business Snapshot 28.2, a swaption can be viewed as a type of bond option.

To give an example of how a swaption might be used, consider a company that knows that in 6 months it will enter into a 5-year floating-rate loan agreement and knows that it will wish to swap the floating interest payments for fixed interest payments to convert the loan into a fixed-rate loan (see Chapter 7 for a discussion of how swaps can be used in this way). At a cost, the company could enter into a swaption giving it the right to receive 6-month LIBOR and pay a certain fixed rate of interest, say 8% per annum, for a 5-year period starting in 6 months. If the fixed rate exchanged for floating on a regular 5-year swap in 6 months turns out to be less than 8% per annum, the company will choose not to exercise the swaption and will enter into a swap agreement in the usual way. However, if it turns out to be greater than 8% per annum, the company will choose to exercise the swaption and will obtain a swap at more favorable terms than those available in the market.

Swaptions, when used in the way just described, provide companies with a guarantee that the fixed rate of interest they will pay on a loan at some future time will not exceed some level. They are an alternative to forward swaps (sometimes called *deferred swaps*). Forward swaps involve no up-front cost but have the disadvantage of obligating the company to enter into a swap agreement. With a swaption, the company is able to benefit from favorable interest rate movements while acquiring protection from unfavorable interest rate movements. The difference between a swaption and a forward swap is analogous to the difference between an option on a foreign currency and a forward contract on the currency.

Valuation of European Swaptions

As explained in Chapter 7 the swap rate for a particular maturity at a particular time is the (mid-market) fixed rate that would be exchanged for LIBOR in a newly issued swap with that maturity. The model usually used to value a European option on a swap assumes that the underlying swap rate at the maturity of the option is lognormal. Consider a swaption where the holder has the right to pay a rate s_K and receive LIBOR on a swap that will last n years starting in T years. We suppose that there are m payments per year under the swap and that the notional principal is L .

Chapter 7 showed that day count conventions may lead to the fixed payments under a swap being slightly different on each payment date. For now we will ignore the effect of

Business Snapshot 28.2 Swaptions and Bond Options

As explained in Chapter 7, an interest rate swap can be regarded as an agreement to exchange a fixed-rate bond for a floating-rate bond. At the start of a swap, the value of the floating-rate bond always equals the principal amount of the swap. A swaption can therefore be regarded as an option to exchange a fixed-rate bond for the principal amount of the swap—that is, a type of bond option.

If a swaption gives the holder the right to pay fixed and receive floating, it is a put option on the fixed-rate bond with strike price equal to the principal. If a swaption gives the holder the right to pay floating and receive fixed, it is a call option on the fixed-rate bond with a strike price equal to the principal.

day count conventions and assume that each fixed payment on the swap is the fixed rate times L/m . The impact of day count conventions is considered at the end of this section.

Suppose that the swap rate for an n -year swap starting at time T proves to be s_T . By comparing the cash flows on a swap where the fixed rate is s_T to the cash flows on a swap where the fixed rate is s_K , it can be seen that the payoff from the swaption consists of a series of cash flows equal to

$$\frac{L}{m} \max(s_T - s_K, 0)$$

The cash flows are received m times per year for the n years of the life of the swap. Suppose that the swap payment dates are T_1, T_2, \dots, T_{mn} , measured in years from today. (It is approximately true that $T_i = T + i/m$.) Each cash flow is the payoff from a call option on s_T with strike price s_K .

Whereas a cap is a portfolio of options on interest rates, a swaption is a single option on the swap rate with repeated payoffs. The standard market model gives the value of a swaption where the holder has the right to pay s_K as

$$\sum_{i=1}^{mn} \frac{L}{m} P(0, T_i) [s_0 N(d_1) - s_K N(d_2)]$$

where

$$d_1 = \frac{\ln(s_0/s_K) + \sigma^2 T/2}{\sigma \sqrt{T}}$$

$$d_2 = \frac{\ln(s_0/s_K) - \sigma^2 T/2}{\sigma \sqrt{T}} = d_1 - \sigma \sqrt{T}$$

s_0 is the forward swap rate at time zero calculated as indicated in equation (27.23), and σ is the volatility of the forward swap rate (so that $\sigma \sqrt{T}$ is the standard deviation of $\ln s_T$).

This is a natural extension of Black's model. The volatility σ is multiplied by \sqrt{T} . The $\sum_{i=1}^{mn} P(0, T_i)$ term is the discount factor for the mn payoffs. Defining A as the value of a contract that pays $1/m$ at times T_i ($1 \leq i \leq mn$), the value of the swaption becomes

$$LA[s_0 N(d_1) - s_K N(d_2)] \quad (28.10)$$

where

$$A = \frac{1}{m} \sum_{i=1}^{mn} P(0, T_i)$$

If the swaption gives the holder the right to receive a fixed rate of s_K instead of paying it, the payoff from the swaption is

$$\frac{L}{m} \max(s_K - s_T, 0)$$

This is a put option on s_T . As before, the payoffs are received at times T_i ($1 \leq i \leq mn$). The standard market model gives the value of the swaption as

$$LA[s_K N(-d_2) - s_0 N(-d_1)] \quad (28.11)$$

Example 28.4

Suppose that the LIBOR yield curve is flat at 6% per annum with continuous compounding. Consider a swaption that gives the holder the right to pay 6.2% in a 3-year swap starting in 5 years. The volatility of the forward swap rate is 20%. Payments are made semiannually and the principal is \$100 million. In this case,

$$A = \frac{1}{2}(e^{-0.06 \times 5.5} + e^{-0.06 \times 6} + e^{-0.06 \times 6.5} + e^{-0.06 \times 7} + e^{-0.06 \times 7.5} + e^{-0.06 \times 8}) = 2.0035$$

A rate of 6% per annum with continuous compounding translates into 6.09% with semiannual compounding. It follows that, in this example, $s_0 = 0.0609$, $s_K = 0.062$, $T = 5$, and $\sigma = 0.2$, so that

$$d_1 = \frac{\ln(0.0609/0.062) + 0.2^2 \times 5/2}{0.2\sqrt{5}} = 0.1836 \quad \text{and} \quad d_2 = d_1 - 0.2\sqrt{5} = -0.2636$$

From equation (28.10), the value of the swaption (in \$ millions) is

$$100 \times 2.0035 \times [0.0609 \times N(0.1836) - 0.062 \times N(-0.2636)] = 2.07$$

or \$2.07. (This is in agreement with the price given by DerivaGem.)

Broker Quotes

Interdealer brokers provide tables of implied volatilities for European swaptions (i.e., values of σ implied by market prices when equations (28.10) and (28.11) are used). The instruments underlying the quotes are usually “at the money” in the sense that the strike swap rate equals the forward swap rate. Table 28.2 shows typical broker quotes

Table 28.2 Typical broker quotes for US European swaptions (mid-market volatilities percent per annum).

Expiration	Swap length (years)						
	1	2	3	4	5	7	10
1 month	17.75	17.75	17.75	17.50	17.00	17.00	16.00
3 months	19.50	19.00	19.00	18.00	17.50	17.00	16.00
6 months	20.00	20.00	19.25	18.50	18.75	17.75	16.75
1 year	22.50	21.75	20.50	20.00	19.50	18.25	16.75
2 years	22.00	22.00	20.75	19.50	19.75	18.25	16.75
3 years	21.50	21.00	20.00	19.25	19.00	17.75	16.50
4 years	20.75	20.25	19.25	18.50	18.25	17.50	16.00
5 years	20.00	19.50	18.50	17.75	17.50	17.00	15.50

provided for the US dollar market. The life of the option is shown on the vertical scale. This varies from 1 month to 5 years. The life of the underlying swap at the maturity of the option is shown on the horizontal scale. This varies from 1 to 10 years. The volatilities in the 1-year column of the table exhibit a hump similar to that discussed for caps earlier. As we move to the columns corresponding to options on longer-lived swaps, the hump persists but it becomes less pronounced.

Theoretical Justification for the Swaption Model

The extension of Black's model used for swaptions can be shown to be internally consistent by considering a world that is forward risk neutral with respect to the annuity A . The analysis in Section 27.4 shows that:

1. The current value of any security is the current value of the annuity multiplied by the expected value of

$$\frac{\text{Security price at time } T}{\text{Value of the annuity at time } T}$$

in this world (see equation (27.25)).

2. The expected value of the swap rate at time T in this world equals the forward swap rate (see equation (27.24)).

The first result shows that the value of the swaption is

$$LAE_A[\max(s_T - s_K, 0)]$$

From the appendix to Chapter 14, this is

$$LA[E_A(s_T)N(d_1) - s_K N(d_2)]$$

where

$$d_1 = \frac{\ln[E_A(s_T)/s_K] + \sigma^2 T/2}{\sigma\sqrt{T}}$$

$$d_2 = \frac{\ln[E_A(s_T)/s_K] - \sigma^2 T/2}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$$

The second result shows that $E_A(s_T)$ equals s_0 . Taken together, the results lead to the swap option pricing formula in equation (28.10). They show that interest rates can be treated as constant for the purposes of discounting provided that the expected swap rate is set equal to the forward swap rate.

The Impact of Day Count Conventions

The above formulas can be made more precise by considering day count conventions. The fixed rate for the swap underlying the swap option is expressed with a day count convention such as actual/365 or 30/360. Suppose that $T_0 = T$ and that, for the applicable day count convention, the accrual fraction corresponding to the time period between T_{i-1} and T_i is a_i . (For example, if T_{i-1} corresponds to March 1 and T_i corresponds to September 1 and the day count is actual/365, $a_i = 184/365 = 0.5041$.)

The formulas that have been presented are then correct with the annuity factor A being defined as

$$A = \sum_{i=1}^{mn} a_i P(0, T_i)$$

As indicated by equation (27.23) the forward swap rate s_0 is given by solving

$$s_0 A = P(0, T) - P(0, T_{mn})$$

28.4 GENERALIZATIONS

We have presented three different versions of Black's model: one for bond options, one for caps, and one for swap options. Each of the models is internally consistent, but they are not consistent with each other. For example, when future bond prices are lognormal, future zero rates and swap rates are not lognormal; when future zero rates are lognormal, future bond prices and swap rates are not lognormal.

The results can be generalized as follows:

1. Consider any instrument that provides a payoff at time T dependent on the value of a bond observed at time T . Its current value is $P(0, T)$ times the expected payoff provided that expectations are calculated in a world where the expected price of the bond equals its forward price.
2. Consider any instrument that provides a payoff at time T^* dependent on the interest rate observed at time T for the period between T and T^* . Its current value is $P(0, T^*)$ times the expected payoff provided that expectations are calculated in a world where the expected value of the underlying interest rate equals the forward interest rate.
3. Consider any instrument that provides a payoff in the form of an annuity. Suppose that the size of the annuity is determined at time T as a function of the n -year swap rate at time T . Suppose also that annuity lasts for n years and payment dates for the annuity are the same as those for the swap. The value of the instrument is A times the expected payoff per year where (a) A is the current value of the annuity when payments are at the rate \$1 per year and (b) expectations are taken in a world where the expected future swap rate equals the forward swap rate.

The first of these results is a generalization of the European bond option model; the second is a generalization of the cap/floor model; the third is a generalization of the swaption model.

28.5 HEDGING INTEREST RATE DERIVATIVES

This section discusses how the material on Greek letters in Chapter 18 can be extended to cover interest rate derivatives.

In the context of interest rate derivatives, delta risk is the risk associated with a shift in the zero curve. Because there are many ways in which the zero curve can shift, many

deltas can be calculated. Some alternatives are:

1. Calculate the impact of a 1-basis-point parallel shift in the zero curve. This is sometimes termed a DV01.
2. Calculate the impact of small changes in the quotes for each of the instruments used to construct the zero curve.
3. Divide the zero curve (or the forward curve) into a number of sections (or buckets). Calculate the impact of shifting the rates in one bucket by 1 basis point, keeping the rest of the initial term structure unchanged. (This is described in Business Snapshot 6.3.)
4. Carry out a principal components analysis as outlined in Section 21.9. Calculate a delta with respect to the changes in each of the first few factors. The first delta then measures the impact of a small, approximately parallel, shift in the zero curve; the second delta measures the impact of a small twist in the zero curve; and so on.

In practice, traders tend to prefer the second approach. They argue that the only way the zero curve can change is if the quote for one of the instruments used to compute the zero curve changes. They therefore feel that it makes sense to focus on the exposures arising from changes in the prices of these instruments.

When several delta measures are calculated, there are many possible gamma measures. Suppose that 10 instruments are used to compute the zero curve and that deltas are calculated by considering the impact of changes in the quotes for each of these. Gamma is a second partial derivative of the form $\partial^2\Pi/\partial x_i \partial x_j$, where Π is the portfolio value. There are 10 choices for x_i and 10 choices for x_j and a total of 55 different gamma measures. This may be “information overload”. One approach is ignore cross-gammas and focus on the 10 partial derivatives where $i = j$. Another is to calculate a single gamma measure as the second partial derivative of the value of the portfolio with respect to a parallel shift in the zero curve. A further possibility is to calculate gammas with respect to the first two factors in a principal components analysis.

The vega of a portfolio of interest rate derivatives measures its exposure to volatility changes. One approach is to calculate the impact on the portfolio of making the same small change to the Black volatilities of all caps and European swap options. However, this assumes that one factor drives all volatilities and may be too simplistic. A better idea is to carry out a principal components analysis on the volatilities of caps and swap options and calculate vega measures corresponding to the first 2 or 3 factors.

SUMMARY

Black’s model and its extensions provide a popular approach for valuing European-style interest rate options. The essence of Black’s model is that the value of the variable underlying the option is assumed to be lognormal at the maturity of the option. In the case of a European bond option, Black’s model assumes that the underlying bond price is lognormal at the option’s maturity. For a cap, the model assumes that the interest rates underlying each of the constituent caplets are lognormally distributed. In the case of a swap option, the model assumes that the underlying swap rate is lognormally distributed. Each of these models is internally consistent, but they are not consistent with each other.

Black's model involves calculating the expected payoff based on the assumption that the expected value of a variable equals its forward value and then discounting the expected payoff at the zero rate observed in the market today. This is the correct procedure for the "plain vanilla" instruments we have considered in this chapter. However, as we shall see in the next chapter, it is not correct in all situations.

FURTHER READING

Black, F., "The Pricing of Commodity Contracts," *Journal of Financial Economics*, 3 (March 1976): 167–79.

Practice Questions (Answers in Solutions Manual)

- 28.1. A company caps 3-month LIBOR at 10% per annum. The principal amount is \$20 million. On a reset date, 3-month LIBOR is 12% per annum. What payment would this lead to under the cap? When would the payment be made?
- 28.2. Explain why a swap option can be regarded as a type of bond option.
- 28.3. Use the Black's model to value a 1-year European put option on a 10-year bond. Assume that the current cash price of the bond is \$125, the strike price is \$110, the 1-year interest rate is 10% per annum, the bond's forward price volatility is 8% per annum, and the present value of the coupons to be paid during the life of the option is \$10.
- 28.4. Explain carefully how you would use (a) spot volatilities and (b) flat volatilities to value a 5-year cap.
- 28.5. Calculate the price of an option that caps the 3-month rate, starting in 15 months' time, at 13% (quoted with quarterly compounding) on a principal amount of \$1,000. The forward interest rate for the period in question is 12% per annum (quoted with quarterly compounding), the 18-month risk-free interest rate (continuously compounded) is 11.5% per annum, and the volatility of the forward rate is 12% per annum.
- 28.6. A bank uses Black's model to price European bond options. Suppose that an implied price volatility for a 5-year option on a bond maturing in 10 years is used to price a 9-year option on the bond. Would you expect the resultant price to be too high or too low? Explain.
- 28.7. Calculate the value of a 4-year European call option on bond that will mature 5 years from today using Black's model. The 5-year cash bond price is \$105, the cash price of a 4-year bond with the same coupon is \$102, the strike price is \$100, the 4-year risk-free interest rate is 10% per annum with continuous compounding, and the volatility for the bond price in 4 years is 2% per annum.
- 28.8. If the yield volatility for a 5-year put option on a bond maturing in 10 years time is specified as 22%, how should the option be valued? Assume that, based on today's interest rates the modified duration of the bond at the maturity of the option will be 4.2 years and the forward yield on the bond is 7%.
- 28.9. What other instrument is the same as a 5-year zero-cost collar where the strike price of the cap equals the strike price of the floor? What does the common strike price equal?
- 28.10. Derive a put–call parity relationship for European bond options.

- 28.11. Derive a put–call parity relationship for European swap options.
- 28.12. Explain why there is an arbitrage opportunity if the implied Black (flat) volatility of a cap is different from that of a floor. Do the broker quotes in Table 28.1 present an arbitrage opportunity?
- 28.13. When a bond's price is lognormal can the bond's yield be negative? Explain your answer.
- 28.14. What is the value of a European swap option that gives the holder the right to enter into a 3-year annual-pay swap in 4 years where a fixed rate of 5% is paid and LIBOR is received? The swap principal is \$10 million. Assume that the yield curve is flat at 5% per annum with annual compounding and the volatility of the swap rate is 20%. Compare your answer with that given by DerivaGem.
- 28.15. Suppose that the yield R on a zero-coupon bond follows the process

$$dR = \mu dt + \sigma dz$$

where μ and σ are functions of R and t , and dz is a Wiener process. Use Itô's lemma to show that the volatility of the zero-coupon bond price declines to zero as it approaches maturity.

- 28.16. Carry out a manual calculation to verify the option prices in Example 28.2.
- 28.17. Suppose that the 1-year, 2-year, 3-year, 4-year, and 5-year zero rates are 6%, 6.4%, 6.7%, 6.9%, and 7%. The price of a 5-year semiannual cap with a principal of \$100 at a cap rate of 8% is \$3. Use DerivaGem to determine:
- The 5-year flat volatility for caps and floors
 - The floor rate in a zero-cost 5-year collar when the cap rate is 8%
- 28.18. Show that $V_1 + f = V_2$, where V_1 is the value of a swaption to pay a fixed rate of s_K and receive LIBOR between times T_1 and T_2 , f is the value of a forward swap to receive a fixed rate of s_K and pay LIBOR between times T_1 and T_2 , and V_2 is the value of a swaption to receive a fixed rate of s_K between times T_1 and T_2 . Deduce that $V_1 = V_2$ when s_K equals the current forward swap rate.
- 28.19. Suppose that zero rates are as in Problem 28.17. Use DerivaGem to determine the value of an option to pay a fixed rate of 6% and receive LIBOR on a 5-year swap starting in 1 year. Assume that the principal is \$100 million, payments are exchanged semiannually, and the swap rate volatility is 21%.
- 28.20. Describe how you would (a) calculate cap flat volatilities from cap spot volatilities and (b) calculate cap spot volatilities from cap flat volatilities.

Further Questions

- 28.21. Consider an 8-month European put option on a Treasury bond that currently has 14.25 years to maturity. The current cash bond price is \$910, the exercise price is \$900, and the volatility for the bond price is 10% per annum. A coupon of \$35 will be paid by the bond in 3 months. The risk-free interest rate is 8% for all maturities up to 1 year. Use Black's model to determine the price of the option. Consider both the case where the strike price corresponds to the cash price of the bond and the case where it corresponds to the quoted price.

- 28.22. Calculate the price of a cap on the 90-day LIBOR rate in 9 months' time when the principal amount is \$1,000. Use Black's model and the following information:
- (a) The quoted 9-month Eurodollar futures price = 92. (Ignore differences between futures and forward rates.)
 - (b) The interest rate volatility implied by a 9-month Eurodollar option = 15% per annum.
 - (c) The current 12-month interest rate with continuous compounding = 7.5% per annum.
 - (d) The cap rate = 8% per annum. (Assume an actual/360 day count.)
- 28.23. Suppose that the LIBOR yield curve is flat at 8% with annual compounding. A swaption gives the holder the right to receive 7.6% in a 5-year swap starting in 4 years. Payments are made annually. The volatility of the forward swap rate is 25% per annum and the principal is \$1 million. Use Black's model to price the swaption. Compare your answer with that given by DerivaGem.
- 28.24. Use the DerivaGem software to value a 5-year collar that guarantees that the maximum and minimum interest rates on a LIBOR-based loan (with quarterly resets) are 7% and 5%, respectively. The LIBOR zero curve (continuously compounded) is currently flat at 6%. Use a flat volatility of 20%. Assume that the principal is \$100.
- 28.25. Use the DerivaGem software to value a European swaption that gives you the right in 2 years to enter into a 5-year swap in which you pay a fixed rate of 6% and receive floating. Cash flows are exchanged semiannually on the swap. The 1-year, 2-year, 5-year, and 10-year zero-coupon interest rates (continuously compounded) are 5%, 6%, 6.5%, and 7%, respectively. Assume a principal of \$100 and a volatility of 15% per annum. Give an example of how the swaption might be used by a corporation. What bond option is equivalent to the swaption?

29

CHAPTER



Convexity, Timing, and Quanto Adjustments

A popular two-step procedure for valuing a European-style derivative is:

1. Calculate the expected payoff by assuming that the expected value of each underlying variable equals its forward value
2. Discount the expected payoff at the risk-free rate applicable for the time period between the valuation date and the payoff date.

We first used this procedure when valuing FRAs and swaps. Chapter 4 shows that an FRA can be valued by calculating the payoff on the assumption that the forward interest rate will be realized and then discounting the payoff at the risk-free rate. Similarly, Chapter 7 extends this, showing that swaps can be valued by calculating cash flows on the assumption that forward rates will be realized and discounting the cash flows at risk-free rates. Chapters 17 and 27 show that Black's model provides a general approach to valuing a wide range of European options—and Black's model is an application of the two-step procedure. The models presented in Chapter 28 for bond options, caps/floors, and swap options are all examples of the two-step procedure.

This raises the issue of whether it is always correct to value European-style interest rate derivatives by using the two-step procedure. The answer is no! For nonstandard interest rate derivatives, it is sometimes necessary to modify the two-step procedure so that an adjustment is made to the forward value of the variable in the first step. This chapter considers three types of adjustments: convexity adjustments, timing adjustments, and quanto adjustments.

29.1 CONVEXITY ADJUSTMENTS

Consider first an instrument that provides a payoff dependent on a bond yield observed at the time of the payoff.

Usually the forward value of a variable S is calculated with reference to a forward contract that pays off $S_T - K$ at time T . It is the value of K that causes the contract to have zero value. As discussed in Section 27.4, forward interest rates and forward yields

are defined differently. A forward interest rate is the rate implied by a forward zero-coupon bond. More generally, a forward bond yield is the yield implied by the forward bond price.

Suppose that B_T is the price of a bond at time T , y_T is its yield, and the (bond pricing) relationship between B_T and y_T is

$$B_T = G(y_T)$$

Define F_0 as the forward bond price at time zero for a transaction maturing at time T and y_0 as the forward bond yield at time zero. The definition of a forward bond yield means that

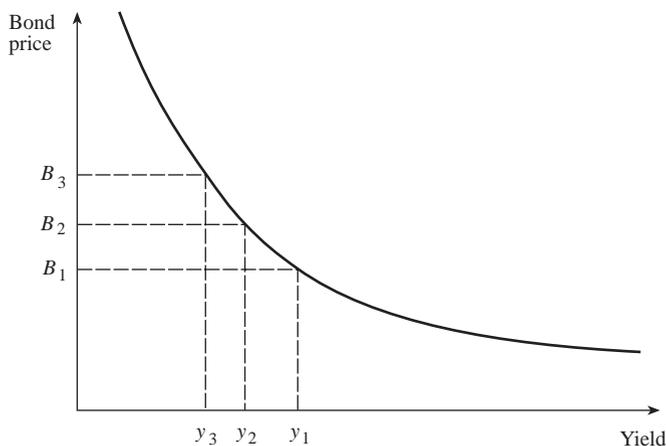
$$F_0 = G(y_0)$$

The function G is nonlinear. This means that, when the expected future bond price equals the forward bond price (so that we are in a world that is forward risk neutral with respect to a zero-coupon bond maturing at time T), the expected future bond yield does not equal the forward bond yield.

This is illustrated in Figure 29.1, which shows the relationship between bond prices and bond yields at time T . For simplicity, suppose that there are only three possible bond prices, B_1 , B_2 , and B_3 and that they are equally likely in a world that is forward risk neutral with respect to $P(t, T)$. Assume that the bond prices are equally spaced, so that $B_2 - B_1 = B_3 - B_2$. The forward bond price is the expected bond price B_2 . The bond prices translate into three equally likely bond yields: y_1 , y_2 , and y_3 . These are not equally spaced. The variable y_2 is the forward bond yield because it is the yield corresponding to the forward bond price. The expected bond yield is the average of y_1 , y_2 , and y_3 and is clearly greater than y_2 .

Consider a derivative that provides a payoff dependent on the bond yield at time T . From equation (27.20), it can be valued by (a) calculating the expected payoff in a world that is forward risk neutral with respect to a zero-coupon bond maturing at time T and (b) discounting at the current risk-free rate for maturity T . We know that the expected bond price equals the forward price in the world being considered. We therefore need to

Figure 29.1 Relationship between bond prices and bond yields at time T .



know the value of the expected bond yield when the expected bond price equals the forward bond price. The analysis in the appendix at the end of this chapter shows that an approximate expression for the required expected bond yield is

$$E_T(y_T) = y_0 - \frac{1}{2}y_0^2\sigma_y^2T \frac{G''(y_0)}{G'(y_0)} \quad (29.1)$$

where G' and G'' denote the first and second partial derivatives of G , E_T denotes expectations in a world that is forward risk neutral with respect to $P(t, T)$, and σ_y is the forward yield volatility. It follows that the expected payoff can be discounted at the current risk-free rate for maturity T provided the expected bond yield is assumed to be

$$y_0 - \frac{1}{2}y_0^2\sigma_y^2T \frac{G''(y_0)}{G'(y_0)}$$

rather than y_0 . The difference between the expected bond yield and the forward bond yield

$$-\frac{1}{2}y_0^2\sigma_y^2T \frac{G''(y_0)}{G'(y_0)}$$

is known as a *convexity adjustment*. It corresponds to the difference between y_2 and the expected yield in Figure 29.1. (The convexity adjustment is positive because $G'(y_0) < 0$ and $G''(y_0) > 0$.)

Application 1: Interest Rates

For a first application of equation (29.1), consider an instrument that provides a cash flow at time T equal to the interest rate between times T and T^* applied to a principal of L . (This example will be useful when we consider LIBOR-in-arrears swaps in Chapter 32.) Note that the interest rate applicable to the time period between times T and T^* is normally paid at time T^* ; here it is assumed that it is paid early, at time T .

The cash flow at time T is $LR_T\tau$, where $\tau = T^* - T$ and R_T is the zero-coupon interest rate applicable to the period between T and T^* (expressed with a compounding period of τ).¹ The variable R_T can be viewed as the yield at time T on a zero-coupon bond maturing at time T^* . The relationship between the price of this bond and its yield is

$$G(y) = \frac{1}{1 + y\tau}$$

From equation (29.1),

$$E_T(R_T) = R_0 - \frac{1}{2}R_0^2\sigma_R^2T \frac{G''(R_0)}{G'(R_0)}$$

or

$$E_T(R_T) = R_0 + \frac{R_0^2\sigma_R^2\tau T}{1 + R_0\tau} \quad (29.2)$$

where R_0 is the forward rate applicable to the period between T and T^* and σ_R is the volatility of the forward rate.

¹ As usual, for ease of exposition we assume actual/actual day counts in our examples.

The value of the instrument is therefore

$$P(0, T)L\tau \left[R_0 + \frac{R_0^2 \sigma_R^2 \tau T}{1 + R_0 \tau} \right]$$

Example 29.1

Consider a derivative that provides a payoff in 3 years equal to the 1-year zero-coupon rate (annually compounded) at that time multiplied by \$1000. Suppose that the zero rate for all maturities is 10% per annum with annual compounding and the volatility of the forward rate applicable to the time period between year 3 and year 4 is 20%. In this case, $R_0 = 0.10$, $\sigma_R = 0.20$, $T = 3$, $\tau = 1$, and $P(0, 3) = 1/1.10^3 = 0.7513$. The value of the derivative is

$$0.7513 \times 1000 \times 1 \times \left[0.10 + \frac{0.10^2 \times 0.20^2 \times 1 \times 3}{1 + 0.10 \times 1} \right]$$

or \$75.95. (This compares with a price of \$75.13 when no convexity adjustment is made.)

Application 2: Swap Rates

Consider next a derivative providing a payoff at time T equal to a swap rate observed at that time. A swap rate is a par yield. For the purposes of calculating a convexity adjustment we can make an approximation and assume that the N -year swap rate at time T equals the yield at that time on an N -year bond with a coupon equal to today's forward swap rate. This enables equation (29.1) to be used.

Example 29.2

Consider an instrument that provides a payoff in 3 years equal to the 3-year swap rate at that time multiplied by \$100. Suppose that payments are made annually on the swap, the zero rate for all maturities is 12% per annum with annual compounding, the volatility for the 3-year forward swap rate in 3 years (implied from swap option prices) is 22%. When the swap rate is approximated as the yield on a 12% bond, the relevant function $G(y)$ is

$$G(y) = \frac{0.12}{1+y} + \frac{0.12}{(1+y)^2} + \frac{1.12}{(1+y)^3}$$

$$G'(y) = -\frac{0.12}{(1+y)^2} - \frac{0.24}{(1+y)^3} - \frac{3.36}{(1+y)^4}$$

$$G''(y) = \frac{0.24}{(1+y)^3} + \frac{0.72}{(1+y)^4} + \frac{13.44}{(1+y)^5}$$

In this case the forward yield y_0 is 0.12, so that $G'(y_0) = -2.4018$ and $G''(y_0) = 8.2546$. From equation (29.1),

$$E_T(y_T) = 0.12 + \frac{1}{2} \times 0.12^2 \times 0.22^2 \times 3 \times \frac{8.2546}{2.4018} = 0.1236$$

A forward swap rate of 0.1236 (= 12.36%) rather than 0.12 should therefore be assumed when valuing the instrument. The instrument is worth

$$\frac{100 \times 0.1236}{1.12^3} = 8.80$$

or \$8.80. (This compares with a price of 8.54 obtained without any convexity adjustment.)

29.2 TIMING ADJUSTMENTS

In this section consider the situation where a market variable V is observed at time T and its value is used to calculate a payoff that occurs at a later time T^* . Define:

V_T : Value of V at time T

$E_T(V_T)$: Expected value of V_T in a world that is forward risk-neutral with respect to $P(t, T)$

$E_{T^*}(V_T)$: Expected value of V_T in a world that is forward risk-neutral with respect to $P(t, T^*)$.

The numeraire ratio when we move from the $P(t, T)$ numeraire to the $P(t, T^*)$ numeraire (see Section 27.8) is

$$W = \frac{P(t, T^*)}{P(t, T)}$$

This is the forward price of a zero-coupon bond lasting between times T and T^* . Define:

σ_V : Volatility of V

σ_W : Volatility of W

ρ_{VW} : Correlation between V and W .

From equation (27.35) the change of numeraire increases the growth rate of V by α_V , where

$$\alpha_V = \rho_{VW}\sigma_V\sigma_W \quad (29.3)$$

This result can be expressed in terms of the forward interest rate between times T and T^* . Define:

R : Forward interest rate for period between T and T^* , expressed with a compounding frequency of m

σ_R : Volatility of R .

The relationship between W and R is

$$W = \frac{1}{(1 + R/m)^{m(T^*-T)}}$$

The relationship between the volatility of W and the volatility of R can be calculated

from Itô's lemma as

$$\sigma_W = -\frac{\sigma_R R(T^* - T)}{1 + R/m}$$

Hence equation (29.3) becomes²

$$\alpha_V = -\frac{\rho_{VR}\sigma_V\sigma_R R(T^* - T)}{1 + R/m}$$

where $\rho_{VR} = -\rho_{VW}$ is the instantaneous correlation between V and R . As an approximation, it can be assumed that R remains constant at its initial value, R_0 , and that the volatilities and correlation in this expression are constant to get, at time zero,

$$E_{T^*}(V_T) = E_T(V_T) \exp\left[-\frac{\rho_{VR}\sigma_V\sigma_R R_0(T^* - T)}{1 + R_0/m} T\right] \quad (29.4)$$

Example 29.3

Consider a derivative that provides a payoff in 6 years equal to the value of a stock index observed in 5 years. Suppose that 1,200 is the forward value of the stock index for a contract maturing in 5 years. Suppose that the volatility of the index is 20%, the volatility of the forward interest rate between years 5 and 6 is 18%, and the correlation between the two is -0.4 . Suppose further that the zero curve is flat at 8% with annual compounding. The results just produced can be used with V defined as the value of the index, $T = 5$, $T^* = 6$, $m = 1$, $R_0 = 0.08$, $\rho_{VR} = -0.4$, $\sigma_V = 0.20$, and $\sigma_R = 0.18$, so that

$$E_{T^*}(V_T) = E_T(V_T) \exp\left[-\frac{-0.4 \times 0.20 \times 0.18 \times 0.08 \times 1}{1 + 0.08} \times 5\right]$$

or $E_{T^*}(V_T) = 1.00535E_T(V_T)$. From the arguments in Chapter 27, $E_T(V_T)$ is the forward price of the index, or 1,200. It follows that $E_{T^*}(V_T) = 1,200 \times 1.00535 = 1206.42$. Using again the arguments in Chapter 27, it follows from equation (27.20) that the value of the derivative is $1206.42 \times P(0, 6)$. In this case, $P(0, 6) = 1/1.08^6 = 0.6302$, so that the value of the derivative is 760.25.

Application 1 Revisited

The analysis just given provides a different way of producing the result in Application 1 of Section 29.1. Using the notation from that application, R_T is the interest rate between T and T^* and R_0 as the forward rate for the period between time T and T^* . From equation (27.22),

$$E_{T^*}(R_T) = R_0$$

Applying equation (29.4) with V equal to R gives

$$E_{T^*}(R_T) = E_T(R_T) \exp\left[-\frac{\sigma_R^2 R_0 \tau}{1 + R_0 \tau} T\right]$$

² Variables R and W are negatively correlated. We can reflect this by setting $\sigma_W = -\sigma_R(T^* - T)/(1 + R/m)$, which is a negative number, and setting $\rho_{VW} = \rho_{VR}$. Alternatively we can change the sign of σ_W so that it is positive and set $\rho_{VW} = -\rho_{VR}$. In either case, we end up with the same formula for α_V .

where $\tau = T^* - T$ (note that $m = 1/\tau$). It follows that

$$R_0 = E_T(R_T) \exp\left[-\frac{\sigma_R^2 R_0 T \tau}{1 + R_0 \tau}\right]$$

or

$$E_T(R_T) = R_0 \exp\left[\frac{\sigma_R^2 R_0 T \tau}{1 + R_0 \tau}\right]$$

Approximating the exponential function gives

$$E_T(R_T) = R_0 + \frac{R_0^2 \sigma_R^2 \tau T}{1 + R_0 \tau}$$

This is the same result as equation (29.2).

29.3 QUANTOS

A *quanto* or *cross-currency derivative* is an instrument where two currencies are involved. The payoff is defined in terms of a variable that is measured in one of the currencies and the payoff is made in the other currency. One example of a quanto is the CME futures contract on the Nikkei discussed in Business Snapshot 5.3. The market variable underlying this contract is the Nikkei 225 index (which is measured in yen), but the contract is settled in US dollars.

Consider a quanto that provides a payoff in currency X at time T . Assume that the payoff depends on the value V of a variable that is observed in currency Y at time T . Define:

$P_X(t, T)$: Value at time t in currency X of a zero-coupon bond paying off 1 unit of currency X at time T

$P_Y(t, T)$: Value at time t in currency Y of a zero-coupon bond paying off 1 unit of currency Y at time T

V_T : Value of V at time T

$E_X(V_T)$: Expected value of V_T in a world that is forward risk neutral with respect to $P_X(t, T)$

$E_Y(V_T)$: Expected value of V_T in a world that is forward risk neutral with respect to $P_Y(t, T)$.

The numeraire ratio when we move from the $P_Y(t, T)$ numeraire to the $P_X(t, T)$ numeraire is

$$W(t) = \frac{P_X(t, T)}{P_Y(t, T)} S(t)$$

where $S(t)$ is the spot exchange rate (units of Y per unit of X) at time t . It follows from this that the numeraire ratio $W(t)$ is the forward exchange rate (units of Y per unit of X) for a contract maturing at time T . Define:

σ_W : Volatility of W

σ_V : Volatility of V

ρ_{VW} : Instantaneous correlation between V and W .

From equation (27.35), the change of numeraire increases the growth rate of V by α_V , where

$$\alpha_V = \rho_{VW}\sigma_V\sigma_W \quad (29.5)$$

If it is assumed that the volatilities and correlation are constant, this means that

$$E_X(V_T) = E_Y(V_T)e^{\rho_{VW}\sigma_V\sigma_W T}$$

or as an approximation

$$E_X(V_T) = E_Y(V_T)(1 + \rho_{VW}\sigma_V\sigma_W T) \quad (29.6)$$

This equation will be used for the valuation of what are known as diff swaps in Chapter 32.

Example 29.4

Suppose that the current value of the Nikkei stock index is 15,000 yen, the 1-year dollar risk-free rate is 5%, the 1-year yen risk-free rate is 2%, and the Nikkei dividend yield is 1%. The forward price of the Nikkei for a 1-year contract denominated in yen can be calculated in the usual way from equation (5.8) as

$$15,000e^{(0.02-0.01)\times 1} = 15,150.75$$

Suppose that the volatility of the index is 20%, the volatility of the 1-year forward yen per dollar exchange rate is 12%, and the correlation between the two is 0.3. In this case $E_Y(V_T) = 15,150.75$, $\sigma_F = 0.20$, $\sigma_W = 0.12$ and $\rho = 0.3$. From equation (29.6), the expected value of the Nikkei in a world that is forward risk neutral with respect to a dollar bond maturing in 1 year is

$$15,150.75e^{0.3\times 0.2\times 0.12\times 1} = 15,260.23$$

This is the forward price of the Nikkei for a contract that provides a payoff in dollars rather than yen. (As an approximation, it is also the futures price of such a contract.)

Using Traditional Risk-Neutral Measures

The forward risk-neutral measure works well when payoffs occur at only one time. In other situations, it is often more appropriate to use the traditional risk-neutral measure. Suppose the process followed by a variable V in the traditional currency- Y risk-neutral world is known and we wish to estimate its process in the traditional currency- X risk-neutral world. Define:

- S : Spot exchange rate (units of Y per unit of X)
- σ_S : Volatility of S
- σ_V : Volatility of V
- ρ : Instantaneous correlation between S and V .

In this case, the change of numeraire is from the money market account in currency Y to the money market account in currency X (with both money market accounts being denominated in currency X). Define g_X as the value of the money market account in currency X and g_Y as the value of the money market account in currency Y . The numeraire ratio is

$$g_X S / g_Y$$

Business Snapshot 29.1 Siegel's Paradox

Consider two currencies, X and Y . Suppose that the interest rates in the two currencies, r_X and r_Y , are constant. Define S as the number of units of currency Y per unit of currency X . As explained in Chapter 5, a currency is an asset that provides a yield at the foreign risk-free rate. The traditional risk-neutral process for S is therefore

$$dS = (r_Y - r_X)S dt + \sigma_S S dz$$

From Itô's lemma, this implies that the process for $1/S$ is

$$d(1/S) = (r_X - r_Y + \sigma_S^2)(1/S) dt - \sigma_S(1/S) dz$$

This leads to what is known as *Siegel's paradox*. Since the expected growth rate of S is $r_Y - r_X$ in a risk-neutral world, symmetry suggests that the expected growth rate of $1/S$ should be $r_X - r_Y$ rather than $r_X - r_Y + \sigma_S^2$.

To understand Siegel's paradox it is necessary to appreciate that the process we have given for S is the risk-neutral process for S in a world where the numeraire is the money market account in currency Y . The process for $1/S$, because it is deduced from the process for S , therefore also assumes that this is the numeraire. Because $1/S$ is the number of units of X per unit of Y , to be symmetrical we should measure the process for $1/S$ in a world where the numeraire is the money market account in currency X . Equation (29.7) shows that when we change the numeraire, from the money market account in currency Y to the money market account in currency X , the growth rate of a variable V increases by $\rho\sigma_V\sigma_S$, where ρ is the correlation between S and V . In this case, $V = 1/S$, so that $\rho = -1$ and $\sigma_V = \sigma_S$. It follows that the change of numeraire causes the growth rate of $1/S$ to increase by $-\sigma_S^2$. This neutralizes the $+\sigma_S^2$ in the process given above for $1/S$. The process for $1/S$ in a world where the numeraire is the money market account in currency X is therefore

$$d(1/S) = (r_X - r_Y)(1/S) dt - \sigma_S(1/S) dz$$

This is symmetrical with the process we started with for S . The paradox has been resolved!

The variables $g_X(t)$ and $g_Y(t)$ have a stochastic drift but zero volatility as explained in Section 27.4. From Itô's lemma it follows that the volatility of the numeraire ratio is σ_S . The change of numeraire therefore involves increasing the expected growth rate of V by

$$\rho\sigma_V\sigma_S \quad (29.7)$$

The market price of risk changes from zero to $\rho\sigma_S$. An application of this result is to Siegel's paradox (see Business Snapshot 29.1).

Example 29.5

A 2-year American option provides a payoff of $S - K$ pounds sterling where S is the level of the S&P 500 at the time of exercise and K is the strike price. The current level of the S&P 500 is 1,200. The risk-free interest rates in sterling and dollars are both constant at 5% and 3%, respectively, the correlation between the dollars/sterling exchange rate and the S&P 500 is 0.2, the volatility of the S&P 500

is 25%, and the volatility of the exchange rate is 12%. The dividend yield on the S&P 500 is 1.5%.

This option can be valued by constructing a binomial tree for the S&P 500 using as the numeraire the money market account in the UK (i.e., using the traditional risk-neutral world as seen from the perspective of a UK investor). From equation (29.7), the change in numeraire from the US to UK money market account leads to an increase in the expected growth rate in the S&P 500 of

$$0.2 \times 0.25 \times 0.12 = 0.006$$

or 0.6%. The growth rate of the S&P 500 using a US dollar numeraire is $3\% - 1.5\% = 1.5\%$. The growth rate using the sterling numeraire is therefore 2.1%. The risk-free interest rate in sterling is 5%. The S&P 500 therefore behaves like an asset providing a dividend yield of $5\% - 2.1\% = 2.9\%$ under the sterling numeraire. Using the parameter values of $S = 1,200$, $K = 1,200$, $r = 0.05$, $q = 0.029$, $\sigma = 0.25$, and $T = 2$ with 100 time steps, DerivaGem estimates the value of the option as £179.83.

SUMMARY

When valuing a derivative providing a payoff at a particular future time it is natural to assume that the variables underlying the derivative equal their forward values and discount at the rate of interest applicable from the valuation date to the payoff date. This chapter has shown that this is not always the correct procedure.

When a payoff depends on a bond yield y observed at time T the expected yield should be assumed to be higher than the forward yield as indicated by equation (29.1). This result can be adapted for situations where a payoff depends on a swap rate. When a variable is observed at time T but the payoff occurs at a later time T^* the forward value of the variable should be adjusted as indicated by equation (29.4). When a variable is observed in one currency but leads to a payoff in another currency the forward value of the variable should also be adjusted. In this case the adjustment is shown in equation (29.6).

These results will be used when nonstandard swaps are considered in Chapter 32.

FURTHER READING

Brotherton-Ratcliffe, R., and B. Iben, "Yield Curve Applications of Swap Products," in *Advanced Strategies in Financial Risk Management* (R. Schwartz and C. Smith, eds.). New York Institute of Finance, 1993.

Jamshidian, F., "Corralling Quantos," *Risk*, March (1994): 71–75.

Reiner, E., "Quanto Mechanics," *Risk*, March (1992), 59–63.

Practice Questions (Answers in Solutions Manual)

- 29.1. Explain how you would value a derivative that pays off $100R$ in 5 years, where R is the 1-year interest rate (annually compounded) observed in 4 years. What difference would it make if the payoff were in (a) 4 years and (b) 6 years?
- 29.2. Explain whether any convexity or timing adjustments are necessary when:
- We wish to value a spread option that pays off every quarter the excess (if any) of the 5-year swap rate over the 3-month LIBOR rate applied to a principal of \$100. The payoff occurs 90 days after the rates are observed.
 - We wish to value a derivative that pays off every quarter the 3-month LIBOR rate minus the 3-month Treasury bill rate. The payoff occurs 90 days after the rates are observed.
- 29.3. Suppose that in Example 28.3 of Section 28.2 the payoff occurs after 1 year (i.e., when the interest rate is observed) rather than in 15 months. What difference does this make to the inputs to Black's model?
- 29.4. The yield curve is flat at 10% per annum with annual compounding. Calculate the value of an instrument where, in 5 years' time, the 2-year swap rate (with annual compounding) is received and a fixed rate of 10% is paid. Both are applied to a notional principal of \$100. Assume that the volatility of the swap rate is 20% per annum. Explain why the value of the instrument is different from zero.
- 29.5. What difference does it make in Problem 29.4 if the swap rate is observed in 5 years, but the exchange of payments takes place in (a) 6 years, and (b) 7 years? Assume that the volatilities of all forward rates are 20%. Assume also that the forward swap rate for the period between years 5 and 7 has a correlation of 0.8 with the forward interest rate between years 5 and 6 and a correlation of 0.95 with the forward interest rate between years 5 and 7.
- 29.6. The price of a bond at time T , measured in terms of its yield, is $G(y_T)$. Assume geometric Brownian motion for the forward bond yield y in a world that is forward risk neutral with respect to a bond maturing at time T . Suppose that the growth rate of the forward bond yield is α and its volatility σ_y .
- Use Itô's lemma to calculate the process for the forward bond price in terms of α , σ_y , y , and $G(y)$.
 - The forward bond price should follow a martingale in the world considered. Use this fact to calculate an expression for α .
 - Show that the expression for α is, to a first approximation, consistent with equation (29.1).
- 29.7. The variable S is an investment asset providing income at rate q measured in currency A. It follows the process

$$dS = \mu_S S dt + \sigma_S S dz$$

in the real world. Defining new variables as necessary, give the process followed by S , and the corresponding market price of risk, in:

- A world that is the traditional risk-neutral world for currency A
- A world that is the traditional risk-neutral world for currency B
- A world that is forward risk neutral with respect to a zero-coupon currency A bond maturing at time T

- (d) A world that is forward risk neutral with respect to a zero-coupon currency B bond maturing at time T .
- 29.8. A call option provides a payoff at time T of $\max(S_T - K, 0)$ yen, where S_T is the dollar price of gold at time T and K is the strike price. Assuming that the storage costs of gold are zero and defining other variables as necessary, calculate the value of the contract.
- 29.9. Suppose that an index of Canadian stocks currently stands at 400. The Canadian dollar is currently worth 0.70 US dollars. The risk-free interest rates in Canada and the US are constant at 6% and 4%, respectively. The dividend yield on the index is 3%. Define Q as the number of Canadian dollars per U.S dollar and S as the value of the index. The volatility of S is 20%, the volatility of Q is 6%, and the correlation between S and Q is 0.4. Use DerivaGem to determine the value of a 2-year American-style call option on the index if:
- It pays off in Canadian dollars the amount by which the index exceeds 400.
 - It pays off in US dollars the amount by which the index exceeds 400.

Further Questions

- 29.10. Consider an instrument that will pay off S dollars in 2 years, where S is the value of the Nikkei index. The index is currently 20,000. The yen/dollar exchange rate is 100 (yen per dollar). The correlation between the exchange rate and the index is 0.3 and the dividend yield on the index is 1% per annum. The volatility of the Nikkei index is 20% and the volatility of the yen/dollar exchange rate is 12%. The interest rates (assumed constant) in the US and Japan are 4% and 2%, respectively.
- What is the value of the instrument?
 - Suppose that the exchange rate at some point during the life of the instrument is Q and the level of the index is S . Show that a US investor can create a portfolio that changes in value by approximately ΔS dollar when the index changes in value by ΔS yen by investing S dollars in the Nikkei and shorting SQ yen.
 - Confirm that this is correct by supposing that the index changes from 20,000 to 20,050 and the exchange rate changes from 100 to 99.7.
 - How would you delta hedge the instrument under consideration?
- 29.11. Suppose that the LIBOR yield curve is flat at 8% (with continuous compounding). The payoff from a derivative occurs in 4 years. It is equal to the 5-year rate minus the 2-year rate at this time, applied to a principal of \$100 with both rates being continuously compounded. (The payoff can be positive or negative.) Calculate the value of the derivative. Assume that the volatility for all rates is 25%. What difference does it make if the payoff occurs in 5 years instead of 4 years? Assume all rates are perfectly correlated.
- 29.12. Suppose that the payoff from a derivative will occur in 10 years and will equal the 3-year US dollar swap rate for a semiannual-pay swap observed at that time applied to a certain principal. Assume that the yield curve is flat at 8% (semiannually compounded) per annum in dollars and 3% (semiannually compounded) in yen. The forward swap rate volatility is 18%, the volatility of the 10-year “yen per dollar” forward exchange rate is 12%, and the correlation between this exchange rate and US dollar interest rates is 0.25.

- (a) What is the value of the derivative if the swap rate is applied to a principal of \$100 million so that the payoff is in dollars?
 - (b) What is its value of the derivative if the swap rate is applied to a principal of 100 million yen so that the payoff is in yen?
- 29.13. The payoff from a derivative will occur in 8 years. It will equal the average of the 1-year interest rates observed at times 5, 6, 7, and 8 years applied to a principal of \$1,000. The yield curve is flat at 6% with annual compounding and the volatilities of all rates are 16%. Assume perfect correlation between all rates. What is the value of the derivative?

APPENDIX

PROOF OF THE CONVEXITY ADJUSTMENT FORMULA

This appendix calculates a convexity adjustment for forward bond yields. Suppose that the payoff from a derivative at time T depends on a bond yield observed at that time. Define:

y_0 : Forward bond yield observed today for a forward contract with maturity T

y_T : Bond yield at time T

B_T : Price of the bond at time T

σ_y : Volatility of the forward bond yield.

Suppose that

$$B_T = G(y_T)$$

Expanding $G(y_T)$ in a Taylor series about $y_T = y_0$ yields the following approximation:

$$B_T = G(y_0) + (y_T - y_0)G'(y_0) + 0.5(y_T - y_0)^2G''(y_0)$$

where G' and G'' are the first and second partial derivatives of G . Taking expectations in a world that is forward risk neutral with respect to a zero-coupon bond maturing at time T gives

$$E_T(B_T) = G(y_0) + E_T(y_T - y_0)G'(y_0) + \frac{1}{2}E_T[(y_T - y_0)^2]G''(y_0)$$

where E_T denotes expectations in this world. The expression $G(y_0)$ is by definition the forward bond price. Also, because of the particular world we are working in, $E_T(B_T)$ equals the forward bond price. Hence $E_T(B_T) = G(y_0)$, so that

$$E_T(y_T - y_0)G'(y_0) + \frac{1}{2}E_T[(y_T - y_0)^2]G''(y_0) = 0$$

The expression $E_T[(y_T - y_0)^2]$ is approximately $\sigma_y^2 y_0^2 T$. Hence it is approximately true that

$$E_T(y_T) = y_0 - \frac{1}{2}y_0^2\sigma_y^2 T \frac{G''(y_0)}{G'(y_0)}$$

This shows that, to obtain the expected bond yield in a world that is forward risk neutral with respect to a zero-coupon bond maturing at time T , the term

$$-\frac{1}{2}y_0^2\sigma_y^2 T \frac{G''(y_0)}{G'(y_0)}$$

should be added to the forward bond yield. This is the result in equation (29.1). For an alternative proof, see Problem 29.6.

30

CHAPTER

Interest Rate Derivatives: Models of the Short Rate



The models for pricing interest rate options that we have presented so far make the assumption that the probability distribution of an interest rate, a bond price, or some other variable at a future point in time is lognormal. They are widely used for valuing instruments such as caps, European bond options, and European swap options. However, they have limitations. They do not provide a description of how interest rates evolve through time. Consequently, they cannot be used for valuing interest rate derivatives that are American-style or structured notes.

This chapter and the next discuss alternative approaches for overcoming these limitations. These involve building what is known as a *term structure model*. This is a model describing the evolution of all zero-coupon interest rates.¹ This chapter focuses on term structure models constructed by specifying the behavior of the short-term interest rate, r .

30.1 BACKGROUND

The short rate, r , at time t is the rate that applies to an infinitesimally short period of time at time t . It is sometimes referred to as the *instantaneous short rate*. Bond prices, option prices, and other derivative prices depend only on the process followed by r in a risk-neutral world. The process for r in the real world is irrelevant. As explained in Chapter 27, the traditional risk-neutral world is a world where, in a very short time period between t and $t + \Delta t$, investors earn on average $r(t) \Delta t$. All processes for r that will be considered in this chapter, except where otherwise stated, are processes in this risk-neutral world.

From equation (27.19), the value at time t of an interest rate derivative that provides a payoff of f_T at time T is

$$\hat{E}[e^{-\bar{r}(T-t)} f_T] \quad (30.1)$$

where \bar{r} is the average value of r in the time interval between t and T , and \hat{E} denotes expected value in the traditional risk-neutral world.

¹ Note that when a term structure model is used we do not need to make the convexity, timing, and quanto adjustments discussed in the previous chapter.

As usual, define $P(t, T)$ as the price at time t of a zero-coupon bond that pays off \$1 at time T . From equation (30.1),

$$P(t, T) = \hat{E}[e^{-\bar{r}(T-t)}] \quad (30.2)$$

If $R(t, T)$ is the continuously compounded interest rate at time t for a term of $T - t$, then

$$P(t, T) = e^{-R(t, T)(T-t)} \quad (30.3)$$

so that

$$R(t, T) = -\frac{1}{T-t} \ln P(t, T) \quad (30.4)$$

and, from equation (30.2),

$$R(t, T) = -\frac{1}{T-t} \ln \hat{E}[e^{-\bar{r}(T-t)}] \quad (30.5)$$

This equation enables the term structure of interest rates at any given time to be obtained from the value of r at that time and the risk-neutral process for r . It shows that, once the process for r has been defined, everything about the initial zero curve and its evolution through time can be determined.

30.2 EQUILIBRIUM MODELS

Equilibrium models usually start with assumptions about economic variables and derive a process for the short rate, r . They then explore what the process for r implies about bond prices and option prices.

In a one-factor equilibrium model, the process for r involves only one source of uncertainty. Usually the risk-neutral process for the short rate is described by an Itô process of the form

$$dr = m(r) dt + s(r) dz$$

The instantaneous drift, m , and instantaneous standard deviation, s , are assumed to be functions of r , but are independent of time. The assumption of a single factor is not as restrictive as it might appear. A one-factor model implies that all rates move in the same direction over any short time interval, but not that they all move by the same amount. The shape of the zero curve can therefore change with the passage of time.

This section considers three one-factor equilibrium models:

$$m(r) = \mu r; \quad s(r) = \sigma r \quad (\text{Rendleman and Bartter model})$$

$$m(r) = a(b - r); \quad s(r) = \sigma \quad (\text{Vasicek model})$$

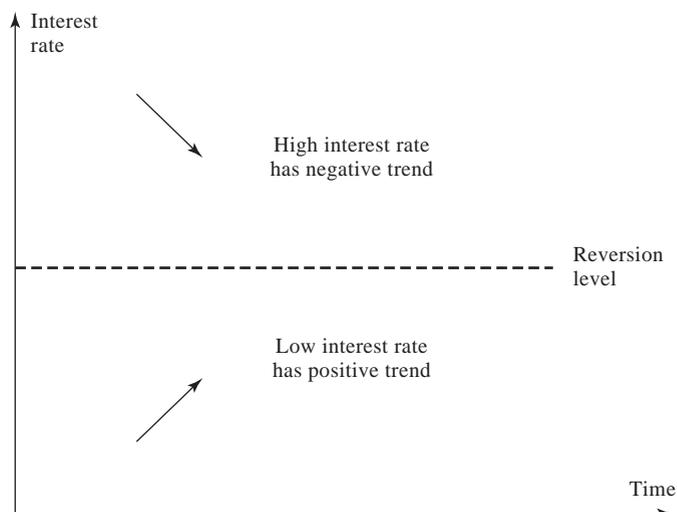
$$m(r) = a(b - r); \quad s(r) = \sigma \sqrt{r} \quad (\text{Cox, Ingersoll, and Ross model})$$

The Rendleman and Bartter Model

In Rendleman and Bartter's model, the risk-neutral process for r is²

$$dr = \mu r dt + \sigma r dz$$

² See R. Rendleman and B. Bartter, "The Pricing of Options on Debt Securities," *Journal of Financial and Quantitative Analysis*, 15 (March 1980): 11–24.

Figure 30.1 Mean reversion.

where μ and σ are constants. This means that r follows geometric Brownian motion. The process for r is of the same type as that assumed for a stock price in Chapter 14. It can be represented using a binomial tree similar to the one used for stocks in Chapter 12.³

The assumption that the short-term interest rate behaves like a stock price is a natural starting point but is less than ideal. One important difference between interest rates and stock prices is that interest rates appear to be pulled back to some long-run average level over time. This phenomenon is known as *mean reversion*. When r is high, mean reversion tends to cause it to have a negative drift; when r is low, mean reversion tends to cause it to have a positive drift. Mean reversion is illustrated in Figure 30.1. The Rendleman and Bartter model does not incorporate mean reversion.

There are compelling economic arguments in favor of mean reversion. When rates are high, the economy tends to slow down and there is low demand for funds from borrowers. As a result, rates decline. When rates are low, there tends to be a high demand for funds on the part of borrowers and rates tend to rise.

The Vasicek Model

In Vasicek's model, the risk-neutral process for r is

$$dr = a(b - r)dt + \sigma dz$$

where a , b , and σ are constants.⁴ This model incorporates mean reversion. The short rate is pulled to a level b at rate a . Superimposed upon this "pull" is a normally distributed stochastic term σdz .

³ The way that the interest rate tree is used is explained later in the chapter.

⁴ See O.A. Vasicek, "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics*, 5 (1977): 177–88.

Vasicek shows that equation (30.2) can be used to obtain the following expression for the price at time t of a zero-coupon bond that pays \$1 at time T :

$$P(t, T) = A(t, T)e^{-B(t, T)r(t)} \quad (30.6)$$

In this equation $r(t)$ is the value of r at time t ,

$$B(t, T) = \frac{1 - e^{-a(T-t)}}{a} \quad (30.7)$$

and

$$A(t, T) = \exp\left[\frac{(B(t, T) - T + t)(a^2b - \sigma^2/2)}{a^2} - \frac{\sigma^2 B(t, T)^2}{4a}\right] \quad (30.8)$$

When $a = 0$, $B(t, T) = T - t$ and $A(t, T) = \exp[\sigma^2(T - t)^3/6]$.

The Cox, Ingersoll, and Ross Model

Cox, Ingersoll, and Ross (CIR) have proposed an alternative model, where⁵

$$dr = a(b - r) dt + \sigma\sqrt{r} dz$$

This has the same mean-reverting drift as Vasicek, but the standard deviation of the change in the short rate in a short period of time is proportional to \sqrt{r} . This means that, as the short-term interest rate increases, the standard deviation increases.

CIR show that, in their model, bond prices have the same general form as those in Vasicek's model,

$$P(t, T) = A(t, T)e^{-B(t, T)r(t)}$$

but the functions $B(t, T)$ and $A(t, T)$ are different:

$$B(t, T) = \frac{2(e^{\gamma(T-t)} - 1)}{(\gamma + a)(e^{\gamma(T-t)} - 1) + 2\gamma}$$

and

$$A(t, T) = \left[\frac{2\gamma e^{(a+\gamma)(T-t)/2}}{(\gamma + a)(e^{\gamma(T-t)} - 1) + 2\gamma}\right]^{2ab/\sigma^2}$$

with $\gamma = \sqrt{a^2 + 2\sigma^2}$.

Properties of Vasicek and CIR

The $A(t, T)$ and $B(t, T)$ functions are different for Vasicek and CIR, but for both models

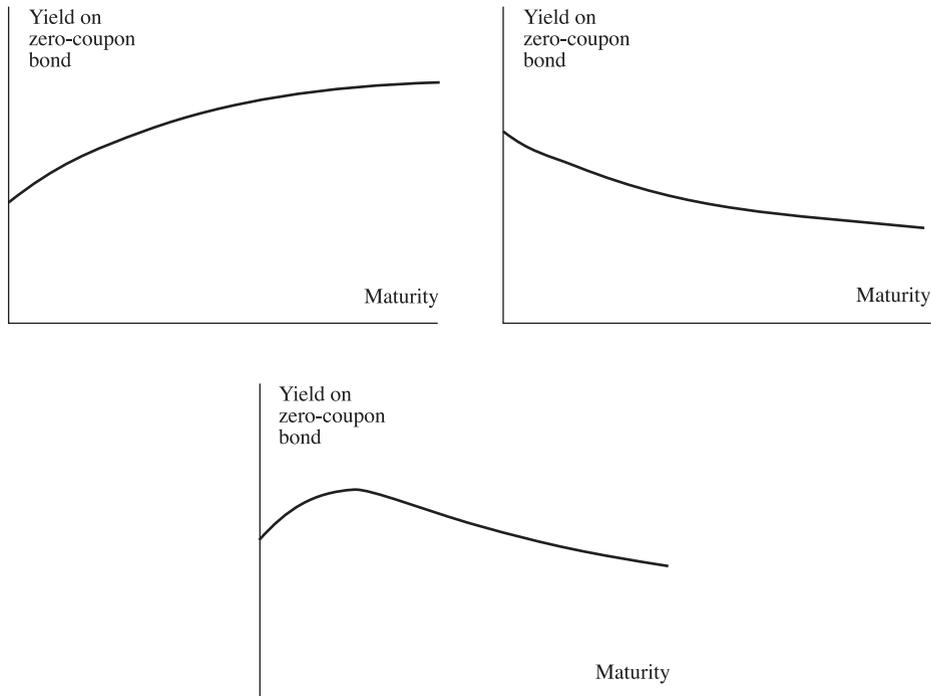
$$P(t, T) = A(t, T)e^{-B(t, T)r(t)}$$

so that

$$\frac{\partial P(t, T)}{r(t)} = -B(t, T)P(t, T) \quad (30.9)$$

⁵ See J.C. Cox, J.E. Ingersoll, and S.A. Ross, "A Theory of the Term Structure of Interest Rates," *Econometrica*, 53 (1985): 385-407.

Figure 30.2 Possible shapes of term structure in the Vasicek and CIR models.



From equation (30.4), the zero rate at time t for a period of $T - t$ is

$$R(t, T) = -\frac{1}{T-t} \ln A(t, T) + \frac{1}{T-t} B(t, T)r(t)$$

This shows that the entire term structure at time t can be determined as a function of $r(t)$ once a , b , and σ have been chosen. The rate $R(t, T)$ is linearly dependent on $r(t)$.⁶ This means that the value of $r(t)$ determines the level of the term structure at time t . The shape of the term structure at time t is independent of $r(t)$, but does depend on t . As shown in Figure 30.2, the shape at a particular time can be upward sloping, downward sloping, or slightly “humped.”

In Chapter 4, we saw that the duration D of a bond or other instrument dependent on interest rates, which has a price of Q , is defined so that

$$\frac{\Delta Q}{Q} = -D\Delta y$$

where ΔQ is the change in Q for a small parallel shift in the yield curve equal to Δy . An

⁶ Some researchers have developed two-factor equilibrium models that give a richer set of possible movements in the term structure than either Vasicek or CIR. See, for example, F. A. Longstaff and E. S. Schwartz, “Interest Rate Volatility and the Term Structure: A Two-Factor General Equilibrium Model,” *Journal of Finance*, 47, 4 (September 1992): 1259–82.

alternative duration measure \hat{D} , which can be used in conjunction with Vasicek or CIR, is

$$\Delta Q/Q = -\hat{D} \Delta r$$

or equivalently

$$\partial Q/\partial r = -\hat{D}Q$$

When Q is the zero-coupon bond, $P(t, T)$, equation (30.9) shows that $\hat{D} = B(t, T)$.

Example 30.1

Consider a zero-coupon bond lasting 4 years. In this case, $D = 4$, so that a 10-basis-point (0.1%) parallel shift in the term structure leads to a decrease of approximately 0.4% in the bond price. If Vasicek's model is used with $a = 0.1$,

$$\hat{D} = B(0, 4) = \frac{(1 - e^{-0.1 \times 4})}{0.1} = 3.29$$

This means that a 10-basis-point increase in the short rate leads to a decrease in the bond price that is approximately 0.329%. The sensitivity of the bond price to movements in the short rate is less than to parallel shifts in the zero curve because of the impact of mean reversion.

When Q is a portfolio of n zero-coupon bonds, $P(t, T_i)$ ($1 \leq i \leq n$), and c_i is the principal of the i th bond, we have

$$\hat{D} = \frac{1}{Q} \frac{\partial Q}{\partial r} = \frac{1}{Q} \sum_{i=1}^n c_i \frac{\partial P(t, T_i)}{\partial r} = \sum_{i=1}^n \frac{c_i P(t, T_i)}{Q} \hat{D}_i$$

where \hat{D}_i is the \hat{D} for $P(t, T_i)$. This shows that the \hat{D} for a coupon-bearing bond can be calculated as a weighted average of the \hat{D} 's for the underlying zero-coupon bonds, similarly to the way the usual duration measure D is calculated (see Table 4.6). A convexity measure for Vasicek and CIR can be defined similarly to the duration measure (see Problem 30.22).

The expected growth rate of $P(t, T)$ in the traditional risk-neutral world at time t is $r(t)$ because $P(t, T)$ is the price of a traded security. Since $P(t, T)$ is a function of $r(t)$, the coefficient of $dz(t)$ in the process for $P(t, T)$ can be calculated from Itô's lemma as $\sigma \partial P(t, T)/\partial r(t)$ for Vasicek and $\sigma\sqrt{r} \partial P(t, T)/\partial r(t)$ for CIR. Substituting from equation (30.9), the processes for $P(t, T)$ in a risk-neutral world are therefore

$$\text{Vasicek: } dP(t, T) = r(t)P(t, T) - \sigma B(t, T)P(t, T) dz(t)$$

$$\text{CIR: } dP(t, T) = r(t)P(t, T) - \sigma\sqrt{r(t)} B(t, T)P(t, T) dz(t)$$

To compare the term structure of interest rates given by Vasicek and CIR for a particular value of r , it makes sense to use the same a and b . However, the Vasicek σ , σ_{vas} , should be chosen to be approximately equal to the CIR σ , σ_{cir} , times \sqrt{r} . For example, if r is 4% and $\sigma_{\text{vas}} = 0.01$, an appropriate value for the σ_{cir} would be $0.01/\sqrt{0.04} = 0.05$. Vasicek gives lower zero-coupon bond yields than CIR. Software for experimenting with the models is at www.rotman.utoronto.ca/~hull/VasicekCIR. Under Vasicek, r can become negative. This is not possible under CIR.⁷

⁷ In CIR, when interest rates get close to zero, the variability of interest rates becomes very small. In all circumstances, negative interest rates are not possible. Zero interest rates are not possible when $2ab \geq \sigma^2$.

Applications of Equilibrium Models

As will be discussed in the next section, when derivatives are being valued it is important that the model used provides an exact fit to the current term structure of interest rates. However, when a Monte Carlo simulation is being carried out over a long period of time for the purposes of scenario analysis, the equilibrium models discussed in this section can be useful tools. A pension fund or insurance company that is interested in the value of its portfolio in 20 years is likely to feel that the precise shape of the current term structure of interest rates has relatively little bearing on its risks.

Once one of the models we have discussed has been chosen, one approach is to determine the parameters from past movements in the short-term interest rate. Data can be collected on daily, weekly, or monthly changes in the short rate and parameters can be determined using linear regression or the maximum-likelihood approach discussed in Section 22.5. Another approach is to use the analytic results to provide as good a fit as possible to the prices of bonds that trade in the market.

There is an important difference between the two approaches. The first approach (fitting historical data) provides parameter estimates in the real world. The second approach (fitting bond prices) provides parameter estimates in the risk-neutral world. When carrying out a scenario analysis, we are interested in modeling the behavior of the short rate in the real world. However, we are also likely to be interested in knowing the complete term structure of interest rates at different times during the life of the Monte Carlo simulation. For this we need risk-neutral parameter estimates.

When we move from the real world to the risk-neutral world, the volatility of the short rate does not change, but the drift does. To determine the change in the drift, it is necessary to make an estimate of the market price of interest rate risk. Ahmad and Wilmott do this by comparing the slope of the zero-coupon yield curve with the real-world drift of the short-term interest rate.⁸ Their estimate of the long-term average market price of interest rate risk for US interest rates is about -1.2 . There is a considerable variation in their estimate of the market price of interest rate risk through time. During stressed market conditions, when the “fear factor” is high (for example, during the 2007–2009 credit crisis), the market price of interest rate risk was found to be a much larger negative number than -1.2 .

Example 30.2

Suppose that the discrete version of Vasicek’s model

$$\Delta r = a(b - r)\Delta t + \sigma\epsilon\sqrt{\Delta t}$$

is used to fit weekly data on a short-term interest rate over a period of 10 years for the purposes of a Monte Carlo simulation. The model parameters can be estimated by regressing Δr on r . Alternatively, maximum-likelihood methods can be used. If r_i is the short-rate at the end of week i ($0 \leq i \leq m$), then the likelihood function is

$$\sum_{i=1}^m \left(-\ln(\sigma^2 \Delta t) - \frac{[r_i - r_{i-1} - a(b - r_{i-1})\Delta t]^2}{\sigma^2 \Delta t} \right)$$

where $\Delta t = 1/52$. Suppose that the best-fit values of a , b , and σ are $a = 0.2$,

⁸ See R. Ahmad and P. Wilmott, “The Market Price of Interest-Rate Risk: Measuring and Modeling Fear and Greed in the Fixed-Income Markets,” *Wilmott*, January 2007, 64–70.

$b = 0.04$, and $\sigma = 0.01$. (These parameters indicate that the short rate reverts to 4.0% with a reversion rate of 20%. The volatility of the short rate at any given time is 1% divided by the short rate.) The short rate can then be simulated in the real world.

To determine the risk-neutral process for r , we note that the proportional drift of r is $a(b - r)/r$ and its volatility is σ/r . From the results in Chapter 27, the proportional drift reduces by $\lambda\sigma/r$ when we move from the real world to the risk-neutral world where λ is the market price of interest rate risk. The process for r in the risk-neutral world is therefore

$$dr = [a(b - r) - \lambda\sigma]dt + \sigma dz$$

or

$$dr = [a(b^* - r)]dt + \sigma dz$$

where

$$b^* = b - \lambda\sigma/a$$

Given the Ahmad and Wilmott results, we might choose to set $\lambda = -1.2$, so that $b^* = 0.04 + 1.2 \times 0.01/0.2 = 0.1$. Equations (30.6) to (30.8) (with $b = b^*$) can then be used to determine the complete term structure of interest rates at any point during the Monte Carlo simulation.

Example 30.3

The Cox–Ingersoll–Ross model

$$dr = a(b - r)dt + \sigma\sqrt{r}dz$$

can be used to value bonds of any maturity using the model's analytic results. Suppose that the values of a , b , and σ that minimize the sum of the squared differences between the market prices of a set of bonds and the prices given by the model are $a = 0.15$, $b = 0.06$, and $\sigma = 0.05$. These values of the parameters give a best-fit risk-neutral process for the short-term interest rate. In this case, the proportional drift in the short rate is $a(b - r)/r$ and the volatility of the short rate σ/\sqrt{r} . From the results in Chapter 27, the proportional drift increases by $\lambda\sigma/\sqrt{r}$ when we move from the risk-neutral world to the real world where λ is the market price of interest rate risk. The real-world process for r is therefore

$$dr = [a(b - r) + \lambda\sigma\sqrt{r}]dt + \sigma\sqrt{r}dz$$

This can be used to simulate the process for the short rate in the real world.⁹ At any given time longer rates can be determined using the risk-neutral process and analytic results. As before, we might choose to set $\lambda = -1.2$.

30.3 NO-ARBITRAGE MODELS

The disadvantage of the equilibrium models we have presented is that they do not automatically fit today's term structure of interest rates. By choosing the parameters judiciously, they can be made to provide an approximate fit to many of the term structures that are encountered in practice. But the fit is not an exact one. Most traders

⁹ In moving between the real world and the risk-neutral world for the Cox–Ingersoll–Ross model, it can be convenient to assume that λ is proportional to \sqrt{r} or $1/\sqrt{r}$, so as to preserve the functional form for the drift.

find this unsatisfactory. Not unreasonably, they argue that they can have very little confidence in the price of a bond option when the model used does not price the underlying bond correctly. A 1% error in the price of the underlying bond may lead to a 25% error in an option price.

A *no-arbitrage model* is a model designed to be exactly consistent with today's term structure of interest rates. The essential difference between an equilibrium and a no-arbitrage model is therefore as follows. In an equilibrium model, today's term structure of interest rates is an output. In a no-arbitrage model, today's term structure of interest rates is an input.

In an equilibrium model, the drift of the short rate (i.e., the coefficient of dt) is not usually a function of time. In a no-arbitrage model, the drift is, in general, dependent on time. This is because the shape of the initial zero curve governs the average path taken by the short rate in the future in a no-arbitrage model. If the zero curve is steeply upward-sloping for maturities between t_1 and t_2 , then r has a positive drift between these times; if it is steeply downward-sloping for these maturities, then r has a negative drift between these times.

It turns out that some equilibrium models can be converted to no-arbitrage models by including a function of time in the drift of the short rate. We now consider the Ho–Lee, Hull–White (one- and two-factor), Black–Derman–Toy, and Black–Karasinski models.

The Ho–Lee Model

Ho and Lee proposed the first no-arbitrage model of the term structure in a paper in 1986.¹⁰ They presented the model in the form of a binomial tree of bond prices with two parameters: the short-rate standard deviation and the market price of risk of the short rate. It has since been shown that the continuous-time limit of the model in a risk-neutral world is

$$dr = \theta(t) dt + \sigma dz \quad (30.10)$$

where σ , the instantaneous standard deviation of the short rate, is constant and $\theta(t)$ is a function of time chosen to ensure that the model fits the initial term structure. The variable $\theta(t)$ defines the average direction that r moves at time t . This is independent of the level of r . Ho and Lee's parameter that concerns the market price of risk is irrelevant when the model is used to price interest rate derivatives.

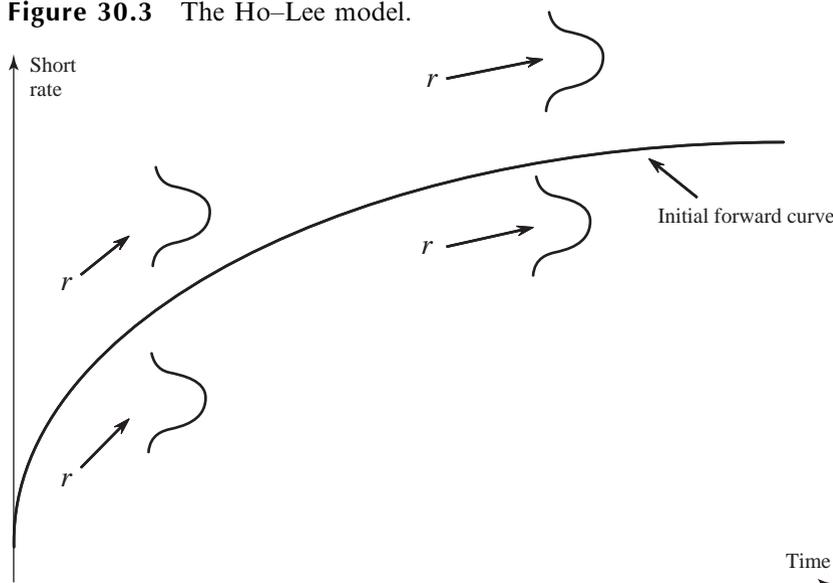
The variable $\theta(t)$ can be calculated analytically (see Problem 30.13). It is

$$\theta(t) = F_t(0, t) + \sigma^2 t \quad (30.11)$$

where the $F(0, t)$ is the instantaneous forward rate for a maturity t as seen at time zero and the subscript t denotes a partial derivative with respect to t . As an approximation, $\theta(t)$ equals $F_t(0, t)$. This means that the average direction that the short rate will be moving in the future is approximately equal to the slope of the instantaneous forward curve. The Ho–Lee model is illustrated in Figure 30.3. Superimposed on the average movement in the short rate is the normally distributed random outcome.

In the Ho–Lee model, zero-coupon bonds and European options on zero-coupon bonds can be valued analytically. The expression for the price of a zero-coupon bond at

¹⁰ See T. S. Y. Ho and S.-B. Lee, "Term Structure Movements and Pricing Interest Rate Contingent Claims," *Journal of Finance*, 41 (December 1986): 1011–29.

Figure 30.3 The Ho–Lee model.

time t in terms of the short rate is

$$P(t, T) = A(t, T)e^{-r(t)(T-t)} \quad (30.12)$$

where

$$\ln A(t, T) = \ln \frac{P(0, T)}{P(0, t)} + (T - t)F(0, t) - \frac{1}{2}\sigma^2 t(T - t)^2$$

In these equations, time zero is today. Times t and T are general times in the future with $T \geq t$. The equations, therefore, define the price of a zero-coupon bond at a future time t in terms of the short rate at time t and the prices of bonds today. The latter can be calculated from today's term structure.

The Hull–White (One-Factor) Model

In a paper published in 1990, Hull and White explored extensions of the Vasicek model that provide an exact fit to the initial term structure.¹¹ One version of the extended Vasicek model that they consider is

$$dr = [\theta(t) - ar]dt + \sigma dz \quad (30.13)$$

or

$$dr = a \left[\frac{\theta(t)}{a} - r \right] dt + \sigma dz$$

where a and σ are constants. This is known as the Hull–White model. It can be characterized as the Ho–Lee model with mean reversion at rate a . Alternatively, it

¹¹ See J. Hull and A. White, "Pricing Interest Rate Derivative Securities," *Review of Financial Studies*, 3, 4 (1990): 573–92.

can be characterized as the Vasicek model with a time-dependent reversion level. At time t , the short rate reverts to $\theta(t)/a$ at rate a . The Ho–Lee model is a particular case of the Hull–White model with $a = 0$.

The model has the same amount of analytic tractability as Ho–Lee. The $\theta(t)$ function can be calculated from the initial term structure (see Problem 30.14):

$$\theta(t) = F_t(0, t) + aF(0, t) + \frac{\sigma^2}{2a}(1 - e^{-2at}) \quad (30.14)$$

The last term in this equation is usually fairly small. If we ignore it, the equation implies that the drift of the process for r at time t is $F_t(0, t) + a[F(0, t) - r]$. This shows that, on average, r follows the slope of the initial instantaneous forward rate curve. When it deviates from that curve, it reverts back to it at rate a . The model is illustrated in Figure 30.4.

Bond prices at time t in the Hull–White model are given by

$$P(t, T) = A(t, T)e^{-B(t, T)r(t)} \quad (30.15)$$

where

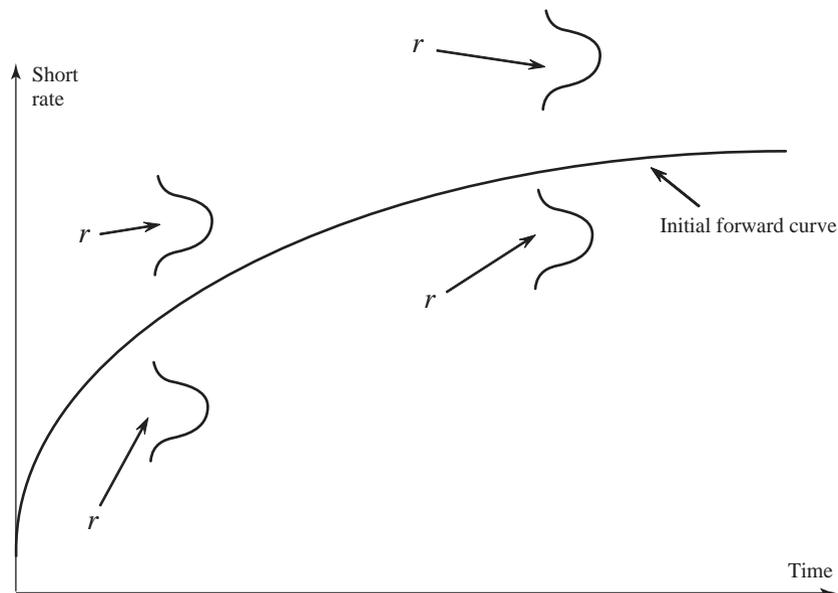
$$B(t, T) = \frac{1 - e^{-a(T-t)}}{a} \quad (30.16)$$

and

$$\ln A(t, T) = \ln \frac{P(0, T)}{P(0, t)} + B(t, T)F(0, t) - \frac{1}{4a^3}\sigma^2(e^{-aT} - e^{-at})^2(e^{2at} - 1) \quad (30.17)$$

Equations (30.15), (30.16), and (30.17) define the price of a zero-coupon bond at a

Figure 30.4 The Hull–White model.



future time t in terms of the short rate at time t and the prices of bonds today. The latter can be calculated from today's term structure.

A method for representing the Hull–White model in the form of a trinomial tree is given later in this chapter. This is useful when American options and other derivatives that cannot be valued analytically are considered.

The Black–Derman–Toy Model

In 1990, Black, Derman, and Toy proposed a binomial-tree model for a lognormal short-rate process.¹² Their procedure for building the binomial tree is explained in Technical Note 23 at www.rotman.utoronto.ca/~hull/TechnicalNotes. It can be shown that the stochastic process corresponding to the model is

$$d \ln r = [\theta(t) - a(t) \ln r] dt + \sigma(t) dz$$

with

$$a(t) = -\frac{\sigma'(t)}{\sigma(t)}$$

where $\sigma'(t)$ is the derivative of σ with respect to t . This model has the advantage over Ho–Lee and Hull–White that the interest rate cannot become negative. The Wiener process dz can cause $\ln(r)$ to be negative, but r itself is always positive. One disadvantage of the model is that there are no analytic properties. A more serious disadvantage is that the way the tree is constructed imposes a relationship between the volatility parameter $\sigma(t)$ and the reversion rate parameter $a(t)$. The reversion rate is positive only if the volatility of the short rate is a decreasing function of time.

In practice, the most useful version of the model is when $\sigma(t)$ is constant. The parameter a is then zero, so that there is no mean reversion and the model reduces to

$$d \ln r = \theta(t) dt + \sigma dz$$

This can be characterized as a lognormal version of the Ho–Lee model.

The Black–Karasinski Model

In 1991, Black and Karasinski developed an extension of the Black–Derman–Toy model where the reversion rate and volatility are determined independently of each other.¹³ The most general version of the model is

$$d \ln r = [\theta(t) - a(t) \ln r] dt + \sigma(t) dz$$

The model is the same as Black–Derman–Toy model except that there is no relation between $a(t)$ and $\sigma(t)$. In practice, $a(t)$ and $\sigma(t)$ are often assumed to be constant, so that the model becomes

$$d \ln r = [\theta(t) - a \ln r] dt + \sigma dz \quad (30.18)$$

As in the case of all the models we are considering, the $\theta(t)$ function is determined to

¹² See F. Black, E. Derman, and W. Toy, "A One-Factor Model of Interest Rates and Its Application to Treasury Bond Prices," *Financial Analysts Journal*, January/February (1990): 33–39.

¹³ See F. Black and P. Karasinski, "Bond and Option Pricing When Short Rates are Lognormal," *Financial Analysts Journal*, July/August (1991): 52–59.

provide an exact fit to the initial term structure of interest rates. The model has no analytic tractability, but later in this chapter we will describe a convenient way of simultaneously determining $\theta(t)$ and representing the process for r in the form of a trinomial tree.

The Hull–White Two-Factor Model

Hull and White have developed a two-factor model:¹⁴

$$df(r) = [\theta(t) + u - af(r)]dt + \sigma_1 dz_1 \quad (30.19)$$

where $f(r)$ is a function of r and u has an initial value of zero and follows the process

$$du = -bu dt + \sigma_2 dz_2$$

As in the one-factor models just considered, the parameter $\theta(t)$ is chosen to make the model consistent with the initial term structure. The stochastic variable u is a component of the reversion level of $f(r)$ and itself reverts to a level of zero at rate b . The parameters a , b , σ_1 , and σ_2 are constants and dz_1 and dz_2 are Wiener processes with instantaneous correlation ρ .

This model provides a richer pattern of term structure movements and a richer pattern of volatilities than one-factor models of r . For more information on the analytical properties of the model and the way a tree can be constructed for it, see Technical Note 14 at www.rotman.utoronto.ca/~hull/TechnicalNotes.

30.4 OPTIONS ON BONDS

Some of the models just presented allow options on zero-coupon bonds to be valued analytically. For the Vasicek, Ho–Lee, and Hull–White models, the price at time zero of a call option that matures at time T on a zero-coupon bond maturing at time s is

$$LP(0, s)N(h) - KP(0, T)N(h - \sigma_p) \quad (30.20)$$

where L is the principal of the bond, K is its strike price, and

$$h = \frac{1}{\sigma_p} \ln \frac{LP(0, s)}{P(0, T)K} + \frac{\sigma_p}{2}$$

The price of a put option on the bond is

$$KP(0, T)N(-h + \sigma_p) - LP(0, s)N(-h)$$

In the case of the Vasicek and Hull–White models,

$$\sigma_p = \frac{\sigma}{a} [1 - e^{-a(s-T)}] \sqrt{\frac{1 - e^{-2aT}}{2a}}$$

¹⁴ See J. Hull and A. White, “Numerical Procedures for Implementing Term Structure Models II: Two-Factor Models,” *Journal of Derivatives*, 2, 2 (Winter 1994): 37–48.

In the case of the Ho–Lee model,

$$\sigma_p = \sigma(s - T)\sqrt{T}$$

Equation (30.20) is essentially the same as Black’s model for pricing bond options in Section 28.1. The forward bond price volatility is σ_p/\sqrt{T} and the standard deviation of the logarithm of the bond price at time T is σ_p . As explained in Section 28.2, an interest rate cap or floor can be expressed as a portfolio of options on zero-coupon bonds. It can, therefore, be valued analytically using the equations just presented.

There are also formulas for valuing options on zero-coupon bonds in the Cox, Ingersoll, and Ross model, which we presented in Section 30.2. These involve integrals of the noncentral chi-square distribution.

Options on Coupon-Bearing Bonds

In a one-factor model of r , all zero-coupon bonds move up in price when r decreases and all zero-coupon bonds move down in price when r increases. As a result, a one-factor model allows a European option on a coupon-bearing bond to be expressed as the sum of European options on zero-coupon bonds. The procedure is as follows:

1. Calculate r^* , the critical value of r for which the price of the coupon-bearing bond equals the strike price of the option on the bond at the option maturity T .
2. Calculate prices of European options with maturity T on the zero-coupon bonds that comprise the coupon-bearing bond. The strike prices of the options equal the values the zero-coupon bonds will have at time T when $r = r^*$.
3. Set the price of the European option on the coupon-bearing bond equal to the sum of the prices on the options on zero-coupon bonds calculated in Step 2.

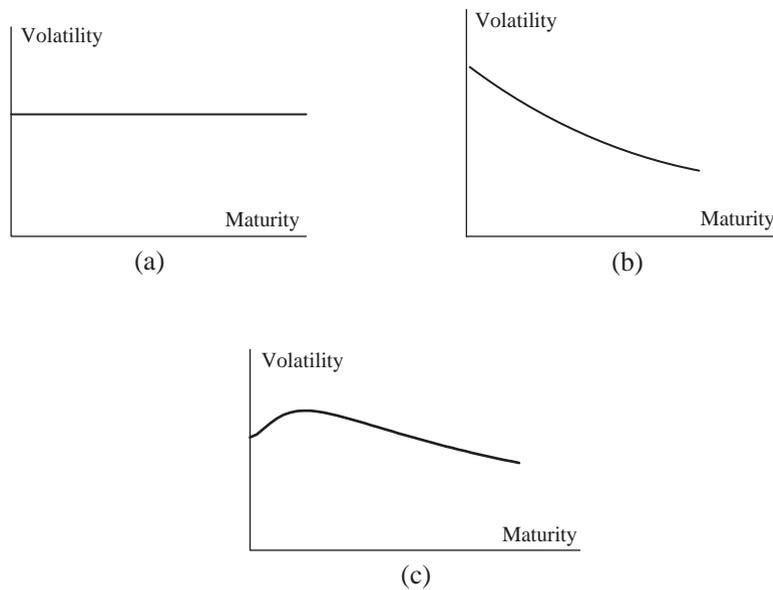
This allows options on coupon-bearing bonds to be valued for the Vasicek, Cox, Ingersoll, and Ross, Ho–Lee, and Hull–White models. As explained in Business Snapshot 28.2, a European swap option can be viewed as an option on a coupon-bearing bond. It can, therefore, be valued using this procedure. For more details on the procedure and a numerical example, see Technical Note 15 at www.rotman.utoronto.ca/~hull/TechnicalNotes.

30.5 VOLATILITY STRUCTURES

The models we have looked at give rise to different volatility environments. Figure 30.5 shows the volatility of the 3-month forward rate as a function of maturity for Ho–Lee, Hull–White one-factor and Hull–White two-factor models. The term structure of interest rates is assumed to be flat.

For Ho–Lee the volatility of the 3-month forward rate is the same for all maturities. In the one-factor Hull–White model the effect of mean reversion is to cause the volatility of the 3-month forward rate to be a declining function of maturity. In the Hull–White two-factor model when parameters are chosen appropriately, the volatility of the 3-month forward rate has a “humped” look. The latter is consistent with empirical evidence and implied cap volatilities discussed in Section 28.2.

Figure 30.5 Volatility of 3-month forward rate as a function of maturity for (a) the Ho–Lee model, (b) the Hull–White one-factor model, and (c) the Hull–White two-factor model (when parameters are chosen appropriately).



30.6 INTEREST RATE TREES

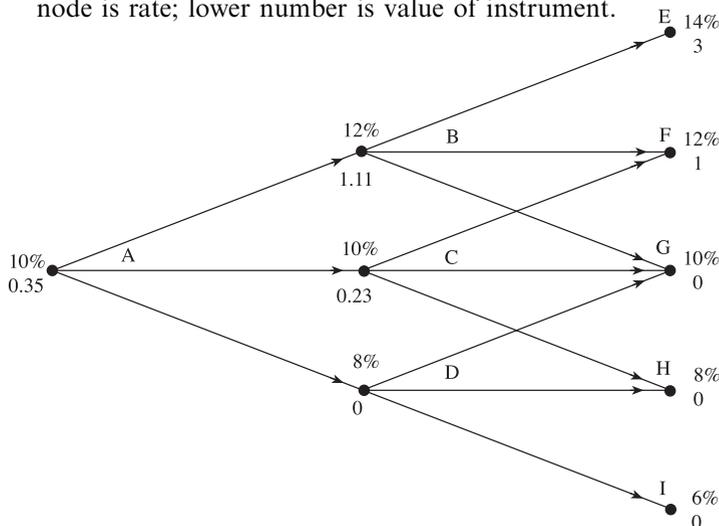
An interest rate tree is a discrete-time representation of the stochastic process for the short rate in much the same way as a stock price tree is a discrete-time representation of the process followed by a stock price. If the time step on the tree is Δt , the rates on the tree are the continuously compounded Δt -period rates. The usual assumption when a tree is constructed is that the Δt -period rate, R , follows the same stochastic process as the instantaneous rate, r , in the corresponding continuous-time model. The main difference between interest rate trees and stock price trees is in the way that discounting is done. In a stock price tree, the discount rate is usually assumed to be the same at each node (or a function of time). In an interest rate tree, the discount rate varies from node to node.

It often proves to be convenient to use a trinomial rather than a binomial tree for interest rates. The main advantage of a trinomial tree is that it provides an extra degree of freedom, making it easier for the tree to represent features of the interest rate process such as mean reversion. As mentioned in Section 20.8, using a trinomial tree is equivalent to using the explicit finite difference method.

Illustration of Use of Trinomial Trees

To illustrate how trinomial interest rate trees are used to value derivatives, consider the simple example shown in Figure 30.6. This is a two-step tree with each time step equal to 1 year in length so that $\Delta t = 1$ year. Assume that the up, middle, and down

Figure 30.6 Example of the use of trinomial interest rate trees. Upper number at each node is rate; lower number is value of instrument.



probabilities are 0.25, 0.50, and 0.25, respectively, at each node. The assumed Δt -period rate is shown as the upper number at each node.¹⁵

The tree is used to value a derivative that provides a payoff at the end of the second time step of

$$\max[100(R - 0.11), 0]$$

where R is the Δt -period rate. The calculated value of this derivative is the lower number at each node. At the final nodes, the value of the derivative equals the payoff. For example, at node E, the value is $100 \times (0.14 - 0.11) = 3$. At earlier nodes, the value of the derivative is calculated using the rollback procedure explained in Chapters 12 and 20. At node B, the 1-year interest rate is 12%. This is used for discounting to obtain the value of the derivative at node B from its values at nodes E, F, and G as

$$[0.25 \times 3 + 0.5 \times 1 + 0.25 \times 0]e^{-0.12 \times 1} = 1.11$$

At node C, the 1-year interest rate is 10%. This is used for discounting to obtain the value of the derivative at node C as

$$(0.25 \times 1 + 0.5 \times 0 + 0.25 \times 0)e^{-0.1 \times 1} = 0.23$$

At the initial node, A, the interest rate is also 10% and the value of the derivative is

$$(0.25 \times 1.11 + 0.5 \times 0.23 + 0.25 \times 0)e^{-0.1 \times 1} = 0.35$$

Nonstandard Branching

It sometimes proves convenient to modify the standard trinomial branching pattern that is used at all nodes in Figure 30.6. Three alternative branching possibilities are shown in

¹⁵ We explain later how the probabilities and rates on an interest rate tree are determined.

Figure 30.7 Alternative branching methods in a trinomial tree.

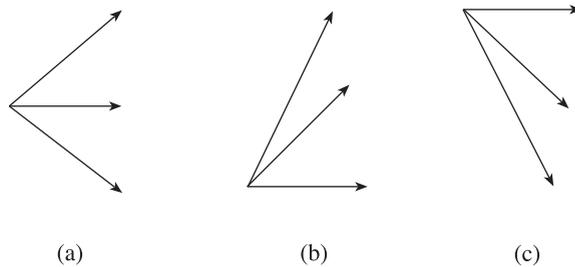


Figure 30.7. The usual branching is shown in Figure 30.7a. It is “up one/straight along/down one”. One alternative to this is “up two/up one/straight along”, as shown in Figure 30.7b. This proves useful for incorporating mean reversion when interest rates are very low. A third branching pattern shown in Figure 30.7c is “straight along/down one/down two”. This is useful for incorporating mean reversion when interest rates are very high. The use of different branching patterns is illustrated in the following section.

30.7 A GENERAL TREE-BUILDING PROCEDURE

Hull and White have proposed a robust two-stage procedure for constructing trinomial trees to represent a wide range of one-factor models.¹⁶ This section first explains how the procedure can be used for the Hull–White model in equation (30.13) and then shows how it can be extended to represent other models, such as Black–Karasinski.

First Stage

The Hull–White model for the instantaneous short rate r is

$$dr = [\theta(t) - ar] dt + \sigma dz$$

We suppose that the time step on the tree is constant and equal to Δt .¹⁷

Assume that the Δt rate, R , follows the same process as r .

$$dR = [\theta(t) - aR] dt + \sigma dz$$

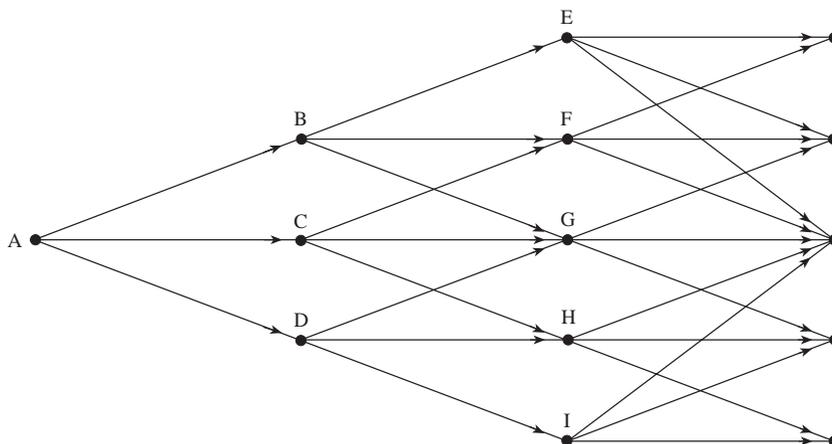
Clearly, this is reasonable in the limit as Δt tends to zero. The first stage in building a tree for this model is to construct a tree for a variable R^* that is initially zero and follows the process

$$dR^* = -aR^* dt + \sigma dz$$

¹⁶ See J. Hull and A. White, “Numerical Procedures for Implementing Term Structure Models I: Single-Factor Models,” *Journal of Derivatives*, 2, 1 (1994): 7–16; and J. Hull and A. White, “Using Hull–White Interest Rate Trees,” *Journal of Derivatives*, (Spring 1996): 26–36.

¹⁷ See Technical Note 16 at www.rotman.utoronto.ca/~hull/TechnicalNotes for a discussion of how nonconstant time steps can be used.

Figure 30.8 Tree for R^* in Hull–White model (first stage).



Node:	A	B	C	D	E	F	G	H	I
R^* (%)	0.000	1.732	0.000	-1.732	3.464	1.732	0.000	-1.732	-3.464
p_u	0.1667	0.1217	0.1667	0.2217	0.8867	0.1217	0.1667	0.2217	0.0867
p_m	0.6666	0.6566	0.6666	0.6566	0.0266	0.6566	0.6666	0.6566	0.0266
p_d	0.1667	0.2217	0.1667	0.1217	0.0867	0.2217	0.1667	0.1217	0.8867

This process is symmetrical about $R^* = 0$. The variable $R^*(t + \Delta t) - R^*(t)$ is normally distributed. If terms of higher order than Δt are ignored, the expected value of $R^*(t + \Delta t) - R^*(t)$ is $-aR^*(t)\Delta t$ and the variance of $R^*(t + \Delta t) - R^*(t)$ is $\sigma^2\Delta t$.

The spacing between interest rates on the tree, ΔR , is set as

$$\Delta R = \sigma\sqrt{3\Delta t}$$

This proves to be a good choice of ΔR from the viewpoint of error minimization.

The objective of the first stage of the procedure is to build a tree similar to that shown in Figure 30.8 for R^* . To do this, it is first necessary to resolve which of the three branching methods shown in Figure 30.7 will apply at each node. This will determine the overall geometry of the tree. Once this is done, the branching probabilities must also be calculated.

Define (i, j) as the node where $t = i\Delta t$ and $R^* = j\Delta R$. (The variable i is a positive integer and j is a positive or negative integer.) The branching method used at a node must lead to the probabilities on all three branches being positive. Most of the time, the branching shown in Figure 30.7a is appropriate. When $a > 0$, it is necessary to switch from the branching in Figure 30.7a to the branching in Figure 30.7c for a sufficiently large j . Similarly, it is necessary to switch from the branching in Figure 30.7a to the branching in Figure 30.7b when j is sufficiently negative. Define j_{\max} as the value of j where we switch from the Figure 30.7a branching to the Figure 30.7c branching and j_{\min} as the value of j where we switch from the Figure 30.7a branching to the Figure 30.7b branching. Hull and White show that probabilities are always positive if j_{\max} is set equal

to the smallest integer greater than $0.184/(a \Delta t)$ and j_{\min} is set equal to $-j_{\max}$.¹⁸ Define p_u , p_m , and p_d as the probabilities of the highest, middle, and lowest branches emanating from the node. The probabilities are chosen to match the expected change and variance of the change in R^* over the next time interval Δt . The probabilities must also sum to unity. This leads to three equations in the three probabilities.

As already mentioned, the mean change in R^* in time Δt is $-aR^* \Delta t$ and the variance of the change is $\sigma^2 \Delta t$. At node (i, j) , $R^* = j \Delta r$. If the branching has the form shown in Figure 30.7a, the p_u , p_m , and p_d at node (i, j) must satisfy the following three equations to match the mean and standard deviation:

$$\begin{aligned} p_u \Delta R - p_d \Delta R &= -aj \Delta R \Delta t \\ p_u \Delta R^2 + p_d \Delta R^2 &= \sigma^2 \Delta t + a^2 j^2 \Delta R^2 \Delta t^2 \\ p_u + p_m + p_d &= 1 \end{aligned}$$

Using $\Delta R = \sigma\sqrt{3\Delta t}$, the solution to these equations is

$$\begin{aligned} p_u &= \frac{1}{6} + \frac{1}{2}(a^2 j^2 \Delta t^2 - aj \Delta t) \\ p_m &= \frac{2}{3} - a^2 j^2 \Delta t^2 \\ p_d &= \frac{1}{6} + \frac{1}{2}(a^2 j^2 \Delta t^2 + aj \Delta t) \end{aligned}$$

Similarly, if the branching has the form shown in Figure 30.7b, the probabilities are

$$\begin{aligned} p_u &= \frac{1}{6} + \frac{1}{2}(a^2 j^2 \Delta t^2 + aj \Delta t) \\ p_m &= -\frac{1}{3} - a^2 j^2 \Delta t^2 - 2aj \Delta t \\ p_d &= \frac{7}{6} + \frac{1}{2}(a^2 j^2 \Delta t^2 + 3aj \Delta t) \end{aligned}$$

Finally, if the branching has the form shown in Figure 30.7c, the probabilities are

$$\begin{aligned} p_u &= \frac{7}{6} + \frac{1}{2}(a^2 j^2 \Delta t^2 - 3aj \Delta t) \\ p_m &= -\frac{1}{3} - a^2 j^2 \Delta t^2 + 2aj \Delta t \\ p_d &= \frac{1}{6} + \frac{1}{2}(a^2 j^2 \Delta t^2 - aj \Delta t) \end{aligned}$$

To illustrate the first stage of the tree construction, suppose that $\sigma = 0.01$, $a = 0.1$, and $\Delta t = 1$ year. In this case, $\Delta R = 0.01\sqrt{3} = 0.0173$, j_{\max} is set equal to the smallest integer greater than $0.184/0.1$, and $j_{\min} = -j_{\max}$. This means that $j_{\max} = 2$ and $j_{\min} = -2$ and the tree is as shown in Figure 30.8. The probabilities on the branches emanating from each node are shown below the tree and are calculated using the equations above for p_u , p_m , and p_d .

Note that the probabilities at each node in Figure 30.8 depend only on j . For example, the probabilities at node B are the same as the probabilities at node F. Furthermore, the tree is symmetrical. The probabilities at node D are the mirror image of the probabilities at node B.

¹⁸ The probabilities are positive for any value of j_{\max} between $0.184/(a \Delta t)$ and $0.816/(a \Delta t)$ and for any value of j_{\min} between $-0.184/(a \Delta t)$ and $-0.816/(a \Delta t)$. Changing the branching at the first possible node proves to be computationally most efficient.

Second Stage

The second stage in the tree construction is to convert the tree for R^* into a tree for R . This is accomplished by displacing the nodes on the R^* -tree so that the initial term structure of interest rates is exactly matched. Define

$$\alpha(t) = R(t) - R^*(t)$$

The $\alpha(t)$'s that apply as the time step Δt on the tree becomes infinitesimally small can be calculated analytically from equation (30.14).¹⁹ However, we want a tree with a finite Δt to match the term structure exactly. We therefore use an iterative procedure to determine the α 's.

Define α_i as $\alpha(i \Delta t)$, the value of R at time $i \Delta t$ on the R -tree minus the corresponding value of R^* at time $i \Delta t$ on the R^* -tree. Define $Q_{i,j}$ as the present value of a security that pays off \$1 if node (i, j) is reached and zero otherwise. The α_i and $Q_{i,j}$ can be calculated using forward induction in such a way that the initial term structure is matched exactly.

Illustration of Second Stage

Suppose that the continuously compounded zero rates in the example in Figure 30.8 are as shown in Table 30.1. The value of $Q_{0,0}$ is 1.0. The value of α_0 is chosen to give the right price for a zero-coupon bond maturing at time Δt . That is, α_0 is set equal to the initial Δt -period interest rate. Because $\Delta t = 1$ in this example, $\alpha_0 = 0.03824$. This defines the position of the initial node on the R -tree in Figure 30.9. The next step is to calculate the values of $Q_{1,1}$, $Q_{1,0}$, and $Q_{1,-1}$. There is a probability of 0.1667 that the $(1, 1)$ node is reached and the discount rate for the first time step is 3.82%. The value of $Q_{1,1}$ is therefore $0.1667e^{-0.0382} = 0.1604$. Similarly, $Q_{1,0} = 0.6417$ and $Q_{1,-1} = 0.1604$.

Once $Q_{1,1}$, $Q_{1,0}$, and $Q_{1,-1}$ have been calculated, α_1 can be determined. It is chosen to give the right price for a zero-coupon bond maturing at time $2\Delta t$. Because $\Delta R = 0.01732$ and $\Delta t = 1$, the price of this bond as seen at node B is $e^{-(\alpha_1 + 0.01732)}$. Similarly, the price as

Table 30.1 Zero rates for example in Figures 30.8 and 30.9.

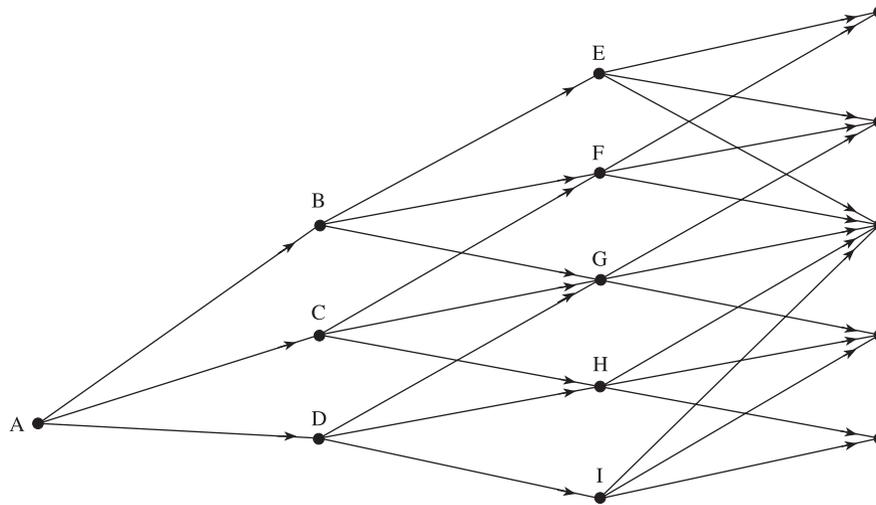
Maturity	Rate (%)
0.5	3.430
1.0	3.824
1.5	4.183
2.0	4.512
2.5	4.812
3.0	5.086

¹⁹ To estimate the instantaneous $\alpha(t)$ analytically, we note that

$$dR = [\theta(t) - aR]dt + \sigma dz \quad \text{and} \quad dR^* = -aR^* dt + \sigma dz$$

so that $d\alpha = [\theta(t) - \alpha(t)]dt$. Using equation (30.14), it can be seen that the solution to this is

$$\alpha(t) = F(0, t) + \frac{\sigma^2}{2a^2}(1 - e^{-at})^2.$$

Figure 30.9 Tree for R in Hull–White model (the second stage).

Node:	A	B	C	D	E	F	G	H	I
R (%)	3.824	6.937	5.205	3.473	9.716	7.984	6.252	4.520	2.788
p_u	0.1667	0.1217	0.1667	0.2217	0.8867	0.1217	0.1667	0.2217	0.0867
p_m	0.6666	0.6566	0.6666	0.6566	0.0266	0.6566	0.6666	0.6566	0.0266
p_d	0.1667	0.2217	0.1667	0.1217	0.0867	0.2217	0.1667	0.1217	0.8867

seen at node C is $e^{-\alpha_1}$ and the price as seen at node D is $e^{-(\alpha_1-0.01732)}$. The price as seen at the initial node A is therefore

$$Q_{1,1}e^{-(\alpha_1+0.01732)} + Q_{1,0}e^{-\alpha_1} + Q_{1,-1}e^{-(\alpha_1-0.01732)} \quad (30.21)$$

From the initial term structure, this bond price should be $e^{-0.04512 \times 2} = 0.9137$. Substituting for the Q 's in equation (30.21),

$$0.1604e^{-(\alpha_1+0.01732)} + 0.6417e^{-\alpha_1} + 0.1604e^{-(\alpha_1-0.01732)} = 0.9137$$

or

$$e^{-\alpha_1}(0.1604e^{-0.01732} + 0.6417 + 0.1604e^{0.01732}) = 0.9137$$

or

$$\alpha_1 = \ln \left[\frac{0.1604e^{-0.01732} + 0.6417 + 0.1604e^{0.01732}}{0.9137} \right] = 0.05205$$

This means that the central node at time Δt in the tree for R corresponds to an interest rate of 5.205% (see Figure 30.9).

The next step is to calculate $Q_{2,2}$, $Q_{2,1}$, $Q_{2,0}$, $Q_{2,-1}$, and $Q_{2,-2}$. The calculations can be shortened by using previously determined Q values. Consider $Q_{2,1}$ as an example. This is the value of a security that pays off \$1 if node F is reached and zero otherwise. Node F can be reached only from nodes B and C. The interest rates at these nodes are 6.937% and 5.205%, respectively. The probabilities associated with the B–F and C–F

branches are 0.6566 and 0.1667. The value at node B of a security that pays \$1 at node F is therefore $0.6566e^{-0.06937}$. The value at node C is $0.1667e^{-0.05205}$. The variable $Q_{2,1}$ is $0.6566e^{-0.06937}$ times the present value of \$1 received at node B plus $0.1667e^{-0.05205}$ times the present value of \$1 received at node C; that is,

$$Q_{2,1} = 0.6566e^{-0.06937} \times 0.1604 + 0.1667e^{-0.05205} \times 0.6417 = 0.1998$$

Similarly, $Q_{2,2} = 0.0182$, $Q_{2,0} = 0.4736$, $Q_{2,-1} = 0.2033$, and $Q_{2,-2} = 0.0189$.

The next step in producing the R -tree in Figure 30.9 is to calculate α_2 . After that, the $Q_{3,j}$'s can then be computed. The variable α_3 can then be calculated, and so on.

Formulas for α 's and Q 's

To express the approach more formally, suppose that the $Q_{i,j}$ have been determined for $i \leq m$ ($m \geq 0$). The next step is to determine α_m so that the tree correctly prices a zero-coupon bond maturing at $(m+1)\Delta t$. The interest rate at node (m, j) is $\alpha_m + j\Delta R$, so that the price of a zero-coupon bond maturing at time $(m+1)\Delta t$ is given by

$$P_{m+1} = \sum_{j=-n_m}^{n_m} Q_{m,j} \exp[-(\alpha_m + j\Delta R)\Delta t] \quad (30.22)$$

where n_m is the number of nodes on each side of the central node at time $m\Delta t$. The solution to this equation is

$$\alpha_m = \frac{\ln \sum_{j=-n_m}^{n_m} Q_{m,j} e^{-j\Delta R\Delta t} - \ln P_{m+1}}{\Delta t}$$

Once α_m has been determined, the $Q_{i,j}$ for $i = m+1$ can be calculated using

$$Q_{m+1,j} = \sum_k Q_{m,k} q(k, j) \exp[-(\alpha_m + k\Delta R)\Delta t]$$

where $q(k, j)$ is the probability of moving from node (m, k) to node $(m+1, j)$ and the summation is taken over all values of k for which this is nonzero.

Extension to Other Models

The procedure that has just been outlined can be extended to more general models of the form

$$df(r) = [\theta(t) - af(r)]dt + \sigma dz \quad (30.23)$$

where f is a monotonic function of r . This family of models has the property that they can fit any term structure.²⁰

²⁰ Not all no-arbitrage models have this property. For example, the extended-CIR model, considered by Cox, Ingersoll, and Ross (1985) and Hull and White (1990), which has the form

$$dr = [\theta(t) - ar]dt + \sigma\sqrt{r}dz$$

cannot fit yield curves where the forward rate declines sharply. This is because the process is not well defined when $\theta(t)$ is negative.

As before, we assume that the Δt period rate, R , follows the same process as r :

$$df(R) = [\theta(t) - af(R)]dt + \sigma dz$$

We start by setting $x = f(R)$, so that

$$dx = [\theta(t) - ax]dt + \sigma dz$$

The first stage is to build a tree for a variable x^* that follows the same process as x except that $\theta(t) = 0$ and the initial value is zero. The procedure here is identical to the procedure already outlined for building a tree such as that in Figure 30.8.

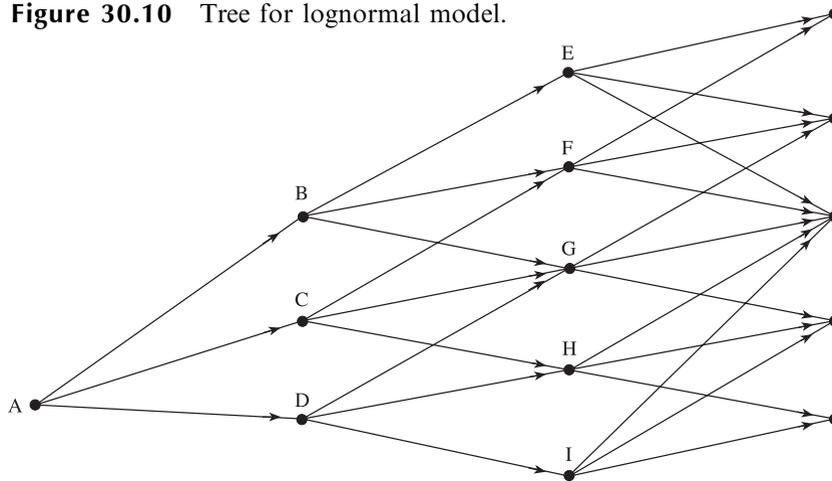
As in Figure 30.9, the nodes at time $i \Delta t$ are then displaced by an amount α_i to provide an exact fit to the initial term structure. The equations for determining α_i and $Q_{i,j}$ inductively are slightly different from those for the $f(R) = R$ case. The value of Q at the first node, $Q_{0,0}$, is set equal to 1. Suppose that the $Q_{i,j}$ have been determined for $i \leq m$ ($m \geq 0$). The next step is to determine α_m so that the tree correctly prices an $(m+1)\Delta t$ zero-coupon bond. Define g as the inverse function of f so that the Δt -period interest rate at the j th node at time $m \Delta t$ is

$$g(\alpha_m + j \Delta x)$$

The price of a zero-coupon bond maturing at time $(m+1)\Delta t$ is given by

$$P_{m+1} = \sum_{j=-n_m}^{n_m} Q_{m,j} \exp[-g(\alpha_m + j \Delta x)\Delta t] \quad (30.24)$$

Figure 30.10 Tree for lognormal model.



Node:	A	B	C	D	E	F	G	H	I
x	-3.373	-2.875	-3.181	-3.487	-2.430	-2.736	-3.042	-3.349	-3.655
R (%)	3.430	5.642	4.154	3.058	8.803	6.481	4.772	3.513	2.587
p_u	0.1667	0.1177	0.1667	0.2277	0.8609	0.1177	0.1667	0.2277	0.0809
p_m	0.6666	0.6546	0.6666	0.6546	0.0582	0.6546	0.6666	0.6546	0.0582
p_d	0.1667	0.2277	0.1667	0.1177	0.0809	0.2277	0.1667	0.1177	0.8609

This equation can be solved using a numerical procedure such as Newton–Raphson. The value α_0 of α when $m = 0$, is $f(R(0))$.

Once α_m has been determined, the $Q_{i,j}$ for $i = m + 1$ can be calculated using

$$Q_{m+1,j} = \sum_k Q_{m,k} q(k, j) \exp[-g(\alpha_m + k \Delta x) \Delta t]$$

where $q(k, j)$ is the probability of moving from node (m, k) to node $(m + 1, j)$ and the summation is taken over all values of k where this is nonzero.

Figure 30.10 shows the results of applying the procedure to the Black–Karasinski model in equation (30.18):

$$d \ln(r) = [\theta(t) - a \ln(r)] dt + \sigma dz$$

when $a = 0.22$, $\sigma = 0.25$, $\Delta t = 0.5$, and the zero rates are as in Table 30.1.

Choosing $f(r)$

Setting $f(r) = r$ leads to the Hull–White model in equation (30.13); setting $f(r) = \ln(r)$ leads to the Black–Karasinski model in equation (30.18). In most circumstances these two models appear to perform about the same in fitting market data on actively traded instruments such as caps and European swap options. The main advantage of the $f(r) = r$ model is its analytic tractability. Its main disadvantage is that negative interest rates are possible. In many circumstances, the probability of negative interest rates occurring under the model is very small, but some analysts are reluctant to use a model where there is any chance at all of negative interest rates. The $f(r) = \ln r$ model has no analytic tractability, but has the advantage that interest rates are always positive. Another advantage is that traders naturally think in terms of σ 's arising from a lognormal model rather than σ 's arising from a normal model.

There is a problem in choosing a satisfactory model for countries with low interest rates. The normal model is unsatisfactory because, when the initial short rate is low, the probability of negative interest rates in the future is no longer negligible. The lognormal model is unsatisfactory because the volatility of rates (i.e., the σ parameter in the lognormal model) is usually much greater when rates are low than when they are high. (For example, a volatility of 100% might be appropriate when the short rate is very low, while 20% might be appropriate when it is 4% or more.) A model that appears to work well is one where $f(r)$ is chosen as a continuous function that is proportional to $\ln r$ when r is very low and proportional to r otherwise.²¹

Using Analytic Results in Conjunction with Trees

When a tree is constructed for the $f(r) = r$ version of the Hull–White model, the analytic results in Section 30.3 can be used to provide the complete term structure and European option prices at each node. It is important to recognize that the interest rate on the tree is the Δt -period rate R . It is not the instantaneous short rate r .

From equations (30.15), (30.16), and (30.17) it can be shown (see Problem 30.21) that

$$P(t, T) = \hat{A}(t, T) e^{-\hat{B}(t, T)R} \quad (30.25)$$

²¹ See J. Hull and A. White “Taking Rates to the Limit,” *Risk*, December (1997): 168–69.

where

$$\ln \hat{A}(t, T) = \ln \frac{P(0, T)}{P(0, t)} - \frac{B(t, T)}{B(t, t + \Delta t)} \ln \frac{P(0, t + \Delta t)}{P(0, t)} - \frac{\sigma^2}{4a} (1 - e^{-2at}) B(t, T) [B(t, T) - B(t, t + \Delta t)] \quad (30.26)$$

and

$$\hat{B}(t, T) = \frac{B(t, T)}{B(t, t + \Delta t)} \Delta t \quad (30.27)$$

(In the case of the Ho–Lee model, we set $\hat{B}(t, T) = T - t$ in these equations.)

Bond prices should therefore be calculated with equation (30.25), and not with equation (30.15).

Example 30.1

Suppose zero rates are as in Table 30.2. The rates for maturities between those indicated are generated using linear interpolation.

Consider a 3-year ($= 3 \times 365$ days) European put option on a zero-coupon bond that will pay 100 in 9 years ($= 9 \times 365$ days). Interest rates are assumed to follow the Hull–White ($f(r) = r$) model. The strike price is 63, $a = 0.1$, and $\sigma = 0.01$. A 3-year tree is constructed and zero-coupon bond prices are calculated analytically at the final nodes as just described. As shown in Table 30.3, the results from the tree are consistent with the analytic price of the option.

This example provides a good test of the implementation of the model because the gradient of the zero curve changes sharply immediately after the expiration of the option. Small errors in the construction and use of the tree are liable to have a big effect on the option values obtained. (The example is used in Sample Application G of the DerivaGem Applications software.)

Table 30.2 Zero curve with all rates continuously compounded, actual/365.

Maturity	Days	Rate (%)
3 days	3	5.01772
1 month	31	4.98284
2 months	62	4.97234
3 months	94	4.96157
6 months	185	4.99058
1 year	367	5.09389
2 years	731	5.79733
3 years	1,096	6.30595
4 years	1,461	6.73464
5 years	1,826	6.94816
6 years	2,194	7.08807
7 years	2,558	7.27527
8 years	2,922	7.30852
9 years	3,287	7.39790
10 years	3,653	7.49015

Table 30.3 Value of a three-year put option on a nine-year zero-coupon bond with a strike price of 63: $a = 0.1$ and $\sigma = 0.01$; zero curve as in Table 30.2.

<i>Steps</i>	<i>Tree</i>	<i>Analytic</i>
10	1.8468	1.8093
30	1.8172	1.8093
50	1.8057	1.8093
100	1.8128	1.8093
200	1.8090	1.8093
500	1.8091	1.8093

Tree for American Bond Options

The DerivaGem software accompanying this book implements the normal and the lognormal model for valuing European and American bond options, caps/floors, and European swap options. Figure 30.11 shows the tree produced by the software when it is used to value a 1.5-year American call option on a 10-year bond using four time steps and the lognormal (Black–Karasinski) model. The parameters used in the lognormal model are $a = 5\%$ and $\sigma = 20\%$. The underlying bond lasts 10 years, has a principal of 100, and pays a coupon of 5% per annum semiannually. The yield curve is flat at 5% per annum. The strike price is 105. As explained in Section 28.1 the strike price can be a cash strike price or a quoted strike price. In this case it is a quoted strike price. The bond price shown on the tree is the cash bond price. The accrued interest at each node is shown below the tree. The cash strike price is calculated as the quoted strike price plus accrued interest. The quoted bond price is the cash bond price minus accrued interest. The payoff from the option is the cash bond price minus the cash strike price. Equivalently it is the quoted bond price minus the quoted strike price.

The tree gives the price of the option as 0.672. A much larger tree with 100 time steps gives the price of the option as 0.703. Note that the price of the 10-year bond cannot be computed analytically when the lognormal model is assumed. It is computed numerically by rolling back through a much larger tree than that shown.

30.8 CALIBRATION

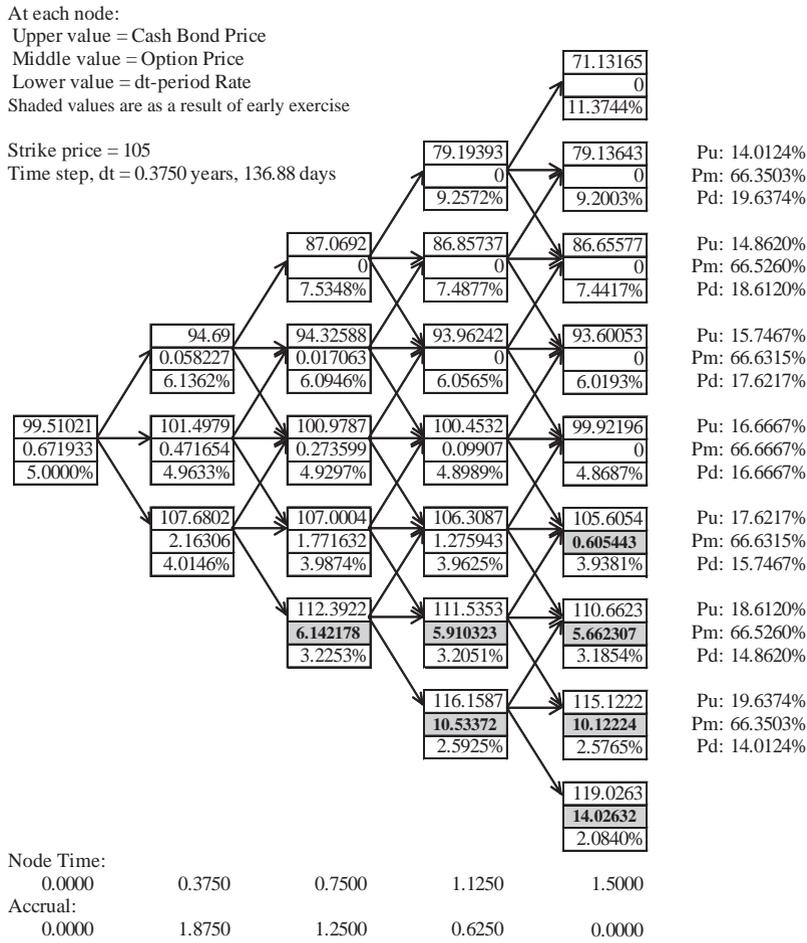
Up to now, we have assumed that the volatility parameters a and σ are known. We now discuss how they are determined. This is known as calibrating the model.

The volatility parameters are determined from market data on actively traded options (e.g., broker quotes on caps and swap options such as those in Tables 28.1 and 28.2). These will be referred to as the *calibrating instruments*. The first stage is to choose a “goodness-of-fit” measure. Suppose there are n calibrating instruments. A popular goodness-of-fit measure is

$$\sum_{i=1}^n (U_i - V_i)^2$$

where U_i is the market price of the i th calibrating instrument and V_i is the price given by

Figure 30.11 Tree, produced by DerivaGem, for valuing an American bond option.



the model for this instrument. The objective of calibration is to choose the model parameters so that this goodness-of-fit measure is minimized.

The number of volatility parameters should not be greater than the number of calibrating instruments. If a and σ are constant, there are only two volatility parameters. The models can be extended so that a or σ , or both, are functions of time. Step functions can be used. Suppose, for example, that a is constant and σ is a function of time. We might choose times t_1, t_2, \dots, t_n and assume $\sigma(t) = \sigma_0$ for $t \leq t_1$, $\sigma(t) = \sigma_i$ for $t_i < t \leq t_{i+1}$ ($1 \leq i \leq n - 1$), and $\sigma(t) = \sigma_n$ for $t > t_n$. There would then be a total of $n + 2$ volatility parameters: $a, \sigma_0, \sigma_1, \dots$, and σ_n .

The minimization of the goodness-of-fit measure can be accomplished using the Levenberg–Marquardt procedure.²² When a or σ , or both, are functions of time, a penalty function is often added to the goodness-of-fit measure so that the functions are

²² For a good description of this procedure, see W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 1988.

“well behaved”. In the example just mentioned, where σ is a step function, an appropriate objective function is

$$\sum_{i=1}^n (U_i - V_i)^2 + \sum_{i=1}^n w_{1,i} (\sigma_i - \sigma_{i-1})^2 + \sum_{i=1}^{n-1} w_{2,i} (\sigma_{i-1} + \sigma_{i+1} - 2\sigma_i)^2$$

The second term provides a penalty for large changes in σ between one step and the next. The third term provides a penalty for high curvature in σ . Appropriate values for $w_{1,i}$ and $w_{2,i}$ are based on experimentation and are chosen to provide a reasonable level of smoothness in the σ function.

The calibrating instruments chosen should be as similar as possible to the instrument being valued. Suppose, for example, that the model is to be used to value a Bermudan-style swap option that lasts 10 years and can be exercised on any payment date between year 5 and year 9 into a swap maturing 10 years from today. The most relevant calibrating instruments are 5×5 , 6×4 , 7×3 , 8×2 , and 9×1 European swap options. (An $n \times m$ European swap option is an n -year option to enter into a swap lasting for m years beyond the maturity of the option.)

The advantage of making a or σ , or both, functions of time is that the models can be fitted more precisely to the prices of instruments that trade actively in the market. The disadvantage is that the volatility structure becomes nonstationary. The volatility term structure given by the model in the future is liable to be quite different from that existing in the market today.²³

A somewhat different approach to calibration is to use all available calibrating instruments to calculate “global-best-fit” a and σ parameters. The parameter a is held fixed at its best-fit value. The model can then be used in the same way as Black–Scholes–Merton. There is a one-to-one relationship between options prices and the σ parameter. The model can be used to convert tables such as Tables 28.1 and 28.2 into tables of implied σ 's.²⁴ These tables can be used to assess the σ most appropriate for pricing the instrument under consideration.

30.9 HEDGING USING A ONE-FACTOR MODEL

Section 28.5 outlined some general approaches to hedging a portfolio of interest rate derivatives. These approaches can be used with the term structure models in this chapter. The calculation of deltas, gammas, and vegas involves making small changes to either the zero curve or the volatility environment and recomputing the value of the portfolio.

Note that, although one factor is often assumed when pricing interest rate derivatives, it is not appropriate to assume only one factor when hedging. For example, the deltas calculated should allow for many different movements in the yield curve, not just those that are possible under the model chosen. The practice of taking account of changes that

²³ For a discussion of the implementation of a model where a and σ are functions of time, see Technical Note 16 at www.rotman.utoronto.ca/~hull/TechnicalNotes.

²⁴ Note that in a term structure model the implied σ 's are not the same as the implied volatilities calculated from Black's model in Tables 28.1 and 28.2. The procedure for computing implied σ 's is as follows. The Black volatilities are converted to prices using Black's model. An iterative procedure is then used to imply the σ parameter in the term structure model from the price.

cannot happen under the model considered, as well as those that can, is known as *outside model hedging* and is standard practice for traders.²⁵ The reality is that relatively simple one-factor models if used carefully usually give reasonable prices for instruments, but good hedging procedures must explicitly or implicitly assume many factors.

SUMMARY

The traditional models of the term structure used in finance are known as equilibrium models. These are useful for understanding potential relationships between variables in the economy, but have the disadvantage that the initial term structure is an output from the model rather than an input to it. When valuing derivatives, it is important that the model used be consistent with the initial term structure observed in the market. No-arbitrage models are designed to have this property. They take the initial term structure as given and define how it can evolve.

This chapter has provided a description of a number of one-factor no-arbitrage models of the short rate. These are robust and can be used in conjunction with any set of initial zero rates. The simplest model is the Ho–Lee model. This has the advantage that it is analytically tractable. Its chief disadvantage is that it implies that all rates are equally variable at all times. The Hull–White model is a version of the Ho–Lee model that includes mean reversion. It allows a richer description of the volatility environment while preserving its analytic tractability. Lognormal one-factor models avoid the possibility of negative interest rates, but have no analytic tractability.

FURTHER READING

Equilibrium Models

- Ahmad, R., and P. Wilmott, “The Market Price of Interest Rate Risk: Measuring and Modeling Fear and Greed in the Fixed Income Markets,” *Wilmott*, January 2007: 64–70.
- Cox, J. C., J. E. Ingersoll, and S. A. Ross, “A Theory of the Term Structure of Interest Rates,” *Econometrica*, 53 (1985): 385–407.
- Longstaff, F. A. and E. S. Schwartz, “Interest Rate Volatility and the Term Structure: A Two Factor General Equilibrium Model,” *Journal of Finance*, 47, 4 (September 1992): 1259–82.
- Vasicek, O. A., “An Equilibrium Characterization of the Term Structure,” *Journal of Financial Economics*, 5 (1977): 177–88.

No-Arbitrage Models

- Black, F., E. Derman, and W. Toy, “A One-Factor Model of Interest Rates and Its Application to Treasury Bond Prices,” *Financial Analysts Journal*, January/February 1990: 33–39.
- Black, F., and P. Karasinski, “Bond and Option Pricing When Short Rates Are Lognormal,” *Financial Analysts Journal*, July/August (1991): 52–59.
- Ho, T. S. Y., and S.-B. Lee, “Term Structure Movements and Pricing Interest Rate Contingent Claims,” *Journal of Finance*, 41 (December 1986): 1011–29.
- Hull, J., and A. White, “Bond Option Pricing Based on a Model for the Evolution of Bond Prices,” *Advances in Futures and Options Research*, 6 (1993): 1–13.

²⁵ A simple example of outside model hedging is in the way that the Black–Scholes–Merton model is used. The Black–Scholes–Merton model assumes that volatility is constant—but traders regularly calculate vega and hedge against volatility changes.

- Hull, J., and A. White, "Pricing Interest Rate Derivative Securities," *The Review of Financial Studies*, 3, 4 (1990): 573–92.
- Hull, J., and A. White, "Using Hull–White Interest Rate Trees," *Journal of Derivatives*, Spring (1996): 26–36.
- Kijima, M., and I. Nagayama, "Efficient Numerical Procedures for the Hull–White Extended Vasicek Model," *Journal of Financial Engineering*, 3 (September/December 1994): 275–92.
- Kijima, M., and I. Nagayama, "A Numerical Procedure for the General One-Factor Interest rate Model" *Journal of Financial Engineering*, 5 (December 1996): 317–37.
- Li, A., P. Ritchken, and L. Sankarasubramanian, "Lattice Models for Pricing American Interest Rate Claims," *Journal of Finance*, 50, 2 (June 1995): 719–37.
- Rebonato, R., *Interest Rate Option Models*. Chichester: Wiley, 1998.

Practice Questions (Answers in Solutions Manual)

- 30.1. What is the difference between an equilibrium model and a no-arbitrage model?
- 30.2. Suppose that the short rate is currently 4% and its standard deviation is 1% per annum. What happens to the standard deviation when the short rate increases to 8% in (a) Vasicek's model; (b) Rendleman and Bartter's model; and (c) the Cox, Ingersoll, and Ross model?
- 30.3. If a stock price were mean reverting or followed a path-dependent process there would be market inefficiency. Why is there not a market inefficiency when the short-term interest rate does so?
- 30.4. Explain the difference between a one-factor and a two-factor interest rate model.
- 30.5. Can the approach described in Section 30.4 for decomposing an option on a coupon-bearing bond into a portfolio of options on zero-coupon bonds be used in conjunction with a two-factor model? Explain your answer.
- 30.6. Suppose that $a = 0.1$ and $b = 0.1$ in both the Vasicek and the Cox, Ingersoll, Ross model. In both models, the initial short rate is 10% and the initial standard deviation of the short-rate change in a short time Δt is $0.02\sqrt{\Delta t}$. Compare the prices given by the models for a zero-coupon bond that matures in year 10.
- 30.7. Suppose that $a = 0.1$, $b = 0.08$, and $\sigma = 0.015$ in Vasicek's model, with the initial value of the short rate being 5%. Calculate the price of a 1-year European call option on a zero-coupon bond with a principal of \$100 that matures in 3 years when the strike price is \$87.
- 30.8. Repeat Problem 30.7 valuing a European put option with a strike of \$87. What is the put–call parity relationship between the prices of European call and put options? Show that the put and call option prices satisfy put–call parity in this case.
- 30.9. Suppose that $a = 0.05$, $b = 0.08$, and $\sigma = 0.015$ in Vasicek's model with the initial short-term interest rate being 6%. Calculate the price of a 2.1-year European call option on a bond that will mature in 3 years. Suppose that the bond pays a coupon of 5% semiannually. The principal of the bond is 100 and the strike price of the option is 99. The strike price is the cash price (not the quoted price) that will be paid for the bond.
- 30.10. Use the answer to Problem 30.9 and put–call parity arguments to calculate the price of a put option that has the same terms as the call option in Problem 30.9.

- 30.11. In the Hull–White model, $a = 0.08$ and $\sigma = 0.01$. Calculate the price of a 1-year European call option on a zero-coupon bond that will mature in 5 years when the term structure is flat at 10%, the principal of the bond is \$100, and the strike price is \$68.
- 30.12. Suppose that $a = 0.05$ and $\sigma = 0.015$ in the Hull–White model with the initial term structure being flat at 6% with semiannual compounding. Calculate the price of a 2.1-year European call option on a bond that will mature in 3 years. Suppose that the bond pays a coupon of 5% per annum semiannually. The principal of the bond is 100 and the strike price of the option is 99. The strike price is the cash price (not the quoted price) that will be paid for the bond.
- 30.13. Use a change of numeraire argument to show that the relationship between the futures rate and forward rate for the Ho–Lee model is as shown in Section 6.3. Use the relationship to verify the expression for $\theta(t)$ given for the Ho–Lee model in equation (30.11). (*Hint*: The futures price is a martingale when the market price of risk is zero. The forward price is a martingale when the market price of risk is a zero-coupon bond maturing at the same time as the forward contract.)
- 30.14. Use a similar approach to that in Problem 30.13 to derive the relationship between the futures rate and the forward rate for the Hull–White model. Use the relationship to verify the expression for $\theta(t)$ given for the Hull–White model in equation (30.14).
- 30.15. Suppose $a = 0.05$, $\sigma = 0.015$, and the term structure is flat at 10%. Construct a trinomial tree for the Hull–White model where there are two time steps, each 1 year in length.
- 30.16. Calculate the price of a 2-year zero-coupon bond from the tree in Figure 30.6.
- 30.17. Calculate the price of a 2-year zero-coupon bond from the tree in Figure 30.9 and verify that it agrees with the initial term structure.
- 30.18. Calculate the price of an 18-month zero-coupon bond from the tree in Figure 30.10 and verify that it agrees with the initial term structure.
- 30.19. What does the calibration of a one-factor term structure model involve?
- 30.20. Use the DerivaGem software to value 1×4 , 2×3 , 3×2 , and 4×1 European swap options to receive fixed and pay floating. Assume that the 1-, 2-, 3-, 4-, and 5-year interest rates are 6%, 5.5%, 6%, 6.5%, and 7%, respectively. The payment frequency on the swap is semiannual and the fixed rate is 6% per annum with semiannual compounding. Use the Hull–White model with $a = 3\%$ and $\sigma = 1\%$. Calculate the volatility implied by Black’s model for each option.
- 30.21. Prove equations (30.25), (30.26), and (30.27).
- 30.22. (a) What is the second partial derivative of $P(t, T)$ with respect to r in the Vasicek and CIR models.
 (b) In Section 30.2, \hat{D} is presented as an alternative to the standard duration measure D . What is a similar alternative \hat{C} to the convexity measure in Section 4.9?
 (c) What is \hat{C} for $P(t, T)$? How would you calculate \hat{C} for a coupon-bearing bond?
 (d) Give a Taylor series expansion for $\Delta P(t, T)$ in terms of Δr and $(\Delta r)^2$ for Vasicek and CIR.
- 30.23. Suppose that short rate r is 4% and its real-world process is

$$dr = 0.1[0.05 - r]dt + 0.01 dz$$

while the risk-neutral process is

$$dr = 0.1[0.11 - r]dt + 0.01 dz$$

- (a) What is the market price of interest rate risk?
- (b) What is the expected return and volatility for a 5-year zero-coupon bond in the risk-neutral world?
- (c) What is the expected return and volatility for the 5-year zero-coupon bond in the real world?

Further Questions

- 30.24. Construct a trinomial tree for the Ho–Lee model where $\sigma = 0.02$. Suppose that the initial zero-coupon interest rate for maturities of 0.5, 1.0, and 1.5 years are 7.5%, 8%, and 8.5%. Use two time steps, each 6 months long. Calculate the value of a zero-coupon bond with a face value of \$100 and a remaining life of 6 months at the ends of the final nodes of the tree. Use the tree to value a 1-year European put option with a strike price of 95 on the bond. Compare the price given by your tree with the analytic price given by DerivaGem.
- 30.25. A trader wishes to compute the price of a 1-year American call option on a 5-year bond with a face value of 100. The bond pays a coupon of 6% semiannually and the (quoted) strike price of the option is \$100. The continuously compounded zero rates for maturities of 6 months, 1 year, 2 years, 3 years, 4 years, and 5 years are 4.5%, 5%, 5.5%, 5.8%, 6.1%, and 6.3%. The best-fit reversion rate for either the normal or the lognormal model has been estimated as 5%.
A 1-year European call option with a (quoted) strike price of 100 on the bond is actively traded. Its market price is \$0.50. The trader decides to use this option for calibration. Use the DerivaGem software with 10 time steps to answer the following questions:
- (a) Assuming a normal model, imply the σ parameter from the price of the European option.
 - (b) Use the σ parameter to calculate the price of the option when it is American.
 - (c) Repeat (a) and (b) for the lognormal model. Show that the model used does not significantly affect the price obtained providing it is calibrated to the known European price.
 - (d) Display the tree for the normal model and calculate the probability of a negative interest rate occurring.
 - (e) Display the tree for the lognormal model and verify that the option price is correctly calculated at the node where, with the notation of Section 30.7, $i = 9$ and $j = -1$.
- 30.26. Use the DerivaGem software to value 1×4 , 2×3 , 3×2 , and 4×1 European swap options to receive floating and pay fixed. Assume that the 1-, 2-, 3-, 3-, and 5-year interest rates are 3%, 3.5%, 3.8%, 4.0%, and 4.1%, respectively. The payment frequency on the swap is semiannual and the fixed rate is 4% per annum with semiannual compounding. Use the lognormal model with $a = 5\%$, $\sigma = 15\%$, and 50 time steps. Calculate the volatility implied by Black's model for each option.
- 30.27. Verify that the DerivaGem software gives Figure 30.11 for the example considered. Use the software to calculate the price of the American bond option for the lognormal and

normal models when the strike price is 95, 100, and 105. In the case of the normal model, assume that $a = 5\%$ and $\sigma = 1\%$. Discuss the results in the context of the heaviness of the tails arguments of Chapter 18.

- 30.28. Modify Sample Application G in the DerivaGem Application Builder software to test the convergence of the price of the trinomial tree when it is used to price a 2-year call option on a 5-year bond with a face value of 100. Suppose that the strike price (quoted) is 100, the coupon rate is 7% with coupons being paid twice a year. Assume that the zero curve is as in Table 30.2. Compare results for the following cases:
- (a) Option is European; normal model with $\sigma = 0.01$ and $a = 0.05$
 - (b) Option is European; lognormal model with $\sigma = 0.15$ and $a = 0.05$
 - (c) Option is American; normal model with $\sigma = 0.01$ and $a = 0.05$
 - (d) Option is American; lognormal model with $\sigma = 0.15$ and $a = 0.05$.
- 30.29. Suppose that the (CIR) process for short-rate movement in the risk-neutral world is

$$dr = a(b - r)dt + \sigma\sqrt{r}dz$$

and the market price of interest rate risk is λ .

- (a) What is the real world process for r ?
- (b) What is the expected return and volatility for a 10-year bond in the risk-neutral world?
- (c) What is the expected return and volatility from a 10-year bond in the real world?



31

CHAPTER

Interest Rate Derivatives: HJM and LMM

The interest rate models discussed in Chapter 30 are widely used for pricing instruments when the simpler models in Chapter 28 are inappropriate. They are easy to implement and, if used carefully, can ensure that most nonstandard interest rate derivatives are priced consistently with actively traded instruments such as interest rate caps, European swap options, and European bond options. Two limitations of the models are:

1. Most involve only one factor (i.e., one source of uncertainty).
2. They do not give the user complete freedom in choosing the volatility structure.

By making the parameters a and σ functions of time, an analyst can use the models so that they fit the volatilities observed in the market today, but as mentioned in Section 30.8 the volatility term structure is then nonstationary. The volatility structure in the future is liable to be quite different from that observed in the market today.

This chapter discusses some general approaches to building term structure models that give the user more flexibility in specifying the volatility environment and allow several factors to be used. The models require much more computation time than the models in Chapter 30. As a result, they are often used for research and development rather than routine pricing.

This chapter also covers the agency mortgage-backed security market in the United States and describes how some of the ideas presented in the chapter can be used to price instruments in that market.

31.1 THE HEATH, JARROW, AND MORTON MODEL

In 1990 David Heath, Bob Jarrow, and Andy Morton (HJM) published an important paper describing the no-arbitrage conditions that must be satisfied by a model of the yield curve.¹ To describe their model, we will use the following notation:

$P(t, T)$: Price at time t of a zero-coupon bond with principal \$1 maturing at time T

¹ See D. Heath, R. A. Jarrow, and A. Morton, "Bond Pricing and the Term Structure of Interest Rates: A New Methodology," *Econometrica*, 60, 1 (1992): 77–105.

- Ω_t : Vector of past and present values of interest rates and bond prices at time t that are relevant for determining bond price volatilities at that time
- $v(t, T, \Omega_t)$: Volatility of $P(t, T)$
- $f(t, T_1, T_2)$: Forward rate as seen at time t for the period between time T_1 and time T_2
- $F(t, T)$: Instantaneous forward rate as seen at time t for a contract maturing at time T
- $r(t)$: Short-term risk-free interest rate at time t
- $dz(t)$: Wiener process driving term structure movements.

Processes for Zero-Coupon Bond Prices and Forward Rates

We start by assuming there is just one factor and will use the traditional risk-neutral world. A zero-coupon bond is a traded security providing no income. Its return in the traditional risk-neutral world must therefore be r . This means that its stochastic process has the form

$$dP(t, T) = r(t)P(t, T) dt + v(t, T, \Omega_t)P(t, T) dz(t) \quad (31.1)$$

As the argument Ω_t indicates, the zero-coupon bond's volatility v can be, in the most general form of the model, any well-behaved function of past and present interest rates and bond prices. Because a bond's price volatility declines to zero at maturity, we must have²

$$v(t, t, \Omega_t) = 0$$

From equation (4.5), the forward rate $f(t, T_1, T_2)$ can be related to zero-coupon bond prices as follows:

$$f(t, T_1, T_2) = \frac{\ln[P(t, T_1)] - \ln[P(t, T_2)]}{T_2 - T_1} \quad (31.2)$$

From equation (31.1) and Itô's lemma,

$$d \ln[P(t, T_1)] = \left[r(t) - \frac{v(t, T_1, \Omega_t)^2}{2} \right] dt + v(t, T_1, \Omega_t) dz(t)$$

and

$$d \ln[P(t, T_2)] = \left[r(t) - \frac{v(t, T_2, \Omega_t)^2}{2} \right] dt + v(t, T_2, \Omega_t) dz(t)$$

so that from equation (31.2)

$$df(t, T_1, T_2) = \frac{v(t, T_2, \Omega_t)^2 - v(t, T_1, \Omega_t)^2}{2(T_2 - T_1)} dt + \frac{v(t, T_1, \Omega_t) - v(t, T_2, \Omega_t)}{T_2 - T_1} dz(t) \quad (31.3)$$

Equation (31.3) shows that the risk-neutral process for f depends solely on the v 's. It depends on r and the P 's only to the extent that the v 's themselves depend on these variables.

² The $v(t, t, \Omega_t) = 0$ condition is equivalent to the assumption that all discount bonds have finite drifts at all times. If the volatility of the bond does not decline to zero at maturity, an infinite drift may be necessary to ensure that the bond's price equals its face value at maturity.

When we put $T_1 = T$ and $T_2 = T + \Delta T$ in equation (31.3) and then take limits as ΔT tends to zero, $f(t, T_1, T_2)$ becomes $F(t, T)$, the coefficient of $dz(t)$ becomes $-v_T(t, T, \Omega_t)$, and the coefficient of dt becomes

$$\frac{1}{2} \frac{\partial [v(t, T, \Omega_t)^2]}{\partial T} = v(t, T, \Omega_t) v_T(t, T, \Omega_t)$$

where the subscript to v denotes a partial derivative. It follows that

$$dF(t, T) = v(t, T, \Omega_t) v_T(t, T, \Omega_t) dt - v_T(t, T, \Omega_t) dz(t) \quad (31.4)$$

Once the function $v(t, T, \Omega_t)$ has been specified, the risk-neutral processes for the $F(t, T)$'s are known.

Equation (31.4) shows that there is a link between the drift and standard deviation of an instantaneous forward rate. This is the key HJM result. Integrating $v_\tau(t, \tau, \Omega_t)$ between $\tau = t$ and $\tau = T$ leads to

$$v(t, T, \Omega_t) - v(t, t, \Omega_t) = \int_t^T v_\tau(t, \tau, \Omega_t) d\tau$$

Because $v(t, t, \Omega_t) = 0$, this becomes

$$v(t, T, \Omega_t) = \int_t^T v_\tau(t, \tau, \Omega_t) d\tau$$

If $m(t, T, \Omega_t)$ and $s(t, T, \Omega_t)$ are the instantaneous drift and standard deviation of $F(t, T)$, so that

$$dF(t, T) = m(t, T, \Omega_t) dt + s(t, T, \Omega_t) dz$$

then it follows from equation (31.4) that

$$m(t, T, \Omega_t) = s(t, T, \Omega_t) \int_t^T s(t, \tau, \Omega_t) d\tau \quad (31.5)$$

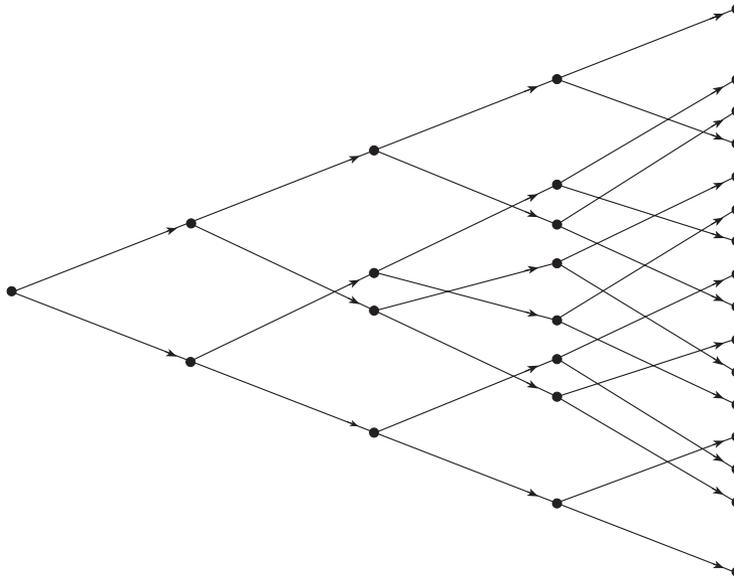
This is the HJM result.

The process for the short rate r in the general HJM model is non-Markov. This means that the process for r at a future time t depends on the path followed by r between now and time t as well as on the value of r at time t .³ This is the key problem in implementing a general HJM model. Monte Carlo simulation has to be used. It is difficult to use a tree to represent term structure movements because the tree is usually nonrecombining. Assuming the model has one factor and the tree is binomial as in Figure 31.1, there are 2^n nodes after n time steps (when $n = 30$, 2^n is about 1 billion).

The HJM model in equation (31.4) is deceptively complex. A particular forward rate $F(t, T)$ is Markov in most applications of the model and can be represented by a recombining tree. However, the same tree cannot be used for all forward rates.

³ For more details, see Technical Note 17 at www.rotman.utoronto.ca/~hull/TechnicalNotes.

Figure 31.1 A nonrecombining tree such as that arising from the general HJM model.



Extension to Several Factors

The HJM result can be extended to the situation where there are several independent factors. Suppose

$$dF(t, T) = m(t, T, \Omega_t) dt + \sum_k s_k(t, T, \Omega_t) dz_k$$

A similar analysis to that just given (see Problem 31.2) shows that

$$m(t, T, \Omega_t) = \sum_k s_k(t, T, \Omega_t) \int_t^T s_k(t, \tau, \Omega_t) d\tau \quad (31.6)$$

31.2 THE LIBOR MARKET MODEL

One drawback of the HJM model is that it is expressed in terms of instantaneous forward rates and these are not directly observable in the market. Another related drawback is that it is difficult to calibrate the model to prices of actively traded instruments. This has led Brace, Gatarek, and Musiela (BGM), Jamshidian, and Miltersen, Sandmann, and Sondermann to propose an alternative.⁴ It is known as the *LIBOR market model* (LMM) or the *BGM model* and it is expressed in terms of the forward rates that traders are used to working with.

⁴ See A. Brace, D. Gatarek, and M. Musiela “The Market Model of Interest Rate Dynamics,” *Mathematical Finance* 7, 2 (1997): 127–55; F. Jamshidian, “LIBOR and Swap Market Models and Measures,” *Finance and Stochastics*, 1 (1997): 293–330; and K. Miltersen, K. Sandmann, and D. Sondermann, “Closed Form Solutions for Term Structure Derivatives with LogNormal Interest Rate,” *Journal of Finance*, 52, 1 (March 1997): 409–30.

The Model

Define $t_0 = 0$ and let t_1, t_2, \dots be the reset times for caps that trade in the market today. In the United States, the most popular caps have quarterly resets, so that it is approximately true that $t_1 = 0.25, t_2 = 0.5, t_3 = 0.75$, and so on. Define $\delta_k = t_{k+1} - t_k$, and

$F_k(t)$: Forward rate between times t_k and t_{k+1} as seen at time t , expressed with a compounding period of δ_k and an actual/actual day count

$m(t)$: Index for the next reset date at time t ; this means that $m(t)$ is the smallest integer such that $t \leq t_{m(t)}$

$\zeta_k(t)$: Volatility of $F_k(t)$ at time t .

Initially, we will assume that there is only one factor.

As shown in Section 27.4, in a world that is forward risk neutral with respect to $P(t, t_{k+1})$, $F_k(t)$ is a martingale and follows the process

$$dF_k(t) = \zeta_k(t)F_k(t) dz \quad (31.7)$$

where dz is a Wiener process.

The process for $P(t, t_k)$ has the form

$$\frac{dP(t, t_k)}{P(t, t_k)} = \dots + v_k(t) dz$$

where $v_k(t)$ is negative because bond prices and interest rates are negatively related.

In practice, it is often most convenient to value interest rate derivatives by working in a world that is always forward risk neutral with respect to a bond maturing at the next reset date. We refer to this as a *rolling forward risk-neutral world*.⁵ In this world we can discount from time t_{k+1} to time t_k using the zero rate observed at time t_k for a maturity t_{k+1} . We do not have to worry about what happens to interest rates between times t_k and t_{k+1} .

At time t the rolling forward risk-neutral world is a world that is forward risk neutral with respect to the bond price, $P(t, t_{m(t)})$. Equation (31.7) gives the process followed by $F_k(t)$ in a world that is forward risk neutral with respect to $P(t, t_{k+1})$. From Section 27.8, it follows that the process followed by $F_k(t)$ in the rolling forward risk-neutral world is

$$dF_k(t) = \zeta_k(t)[v_{m(t)}(t) - v_{k+1}(t)]F_k(t) dt + \zeta_k(t)F_k(t) dz \quad (31.8)$$

The relationship between forward rates and bond prices is

$$\frac{P(t, t_i)}{P(t, t_{i+1})} = 1 + \delta_i F_i(t)$$

or

$$\ln P(t, t_i) - \ln P(t, t_{i+1}) = \ln[1 + \delta_i F_i(t)]$$

Itô's lemma can be used to calculate the process followed by both the left-hand side and

⁵ In the terminology of Section 27.4, this world corresponds to using a "rolling CD" as the numeraire. A rolling CD (certificate of deposit) is one where we start with \$1, buy a bond maturing at time t_1 , reinvest the proceeds at time t_1 in a bond maturing at time t_2 , reinvest the proceeds at time t_2 in a bond maturing at time t_3 , and so on. (Strictly speaking, the interest rate trees we constructed in Chapter 30 are in a rolling forward risk-neutral world rather than the traditional risk-neutral world.) The numeraire is a CD rolled over at the end of each time step.

the right-hand side of this equation. Equating the coefficients of dz gives⁶

$$v_i(t) - v_{i+1}(t) = \frac{\delta_i F_i(t) \zeta_i(t)}{1 + \delta_i F_i(t)} \quad (31.9)$$

so that from equation (31.8) the process followed by $F_k(t)$ in the rolling forward risk-neutral world is

$$\frac{dF_k(t)}{F_k(t)} = \sum_{i=m(t)}^k \frac{\delta_i F_i(t) \zeta_i(t) \zeta_k(t)}{1 + \delta_i F_i(t)} dt + \zeta_k(t) dz \quad (31.10)$$

The HJM result in equation (31.4) is the limiting case of this as the δ_i tend to zero (see Problem 31.7).

Forward Rate Volatilities

The model can be simplified by assuming that $\zeta_k(t)$ is a function only of the number of whole accrual periods between the next reset date and time t_k . Define Λ_i as the value of $\zeta_k(t)$ when there are i such accrual periods. This means that $\zeta_k(t) = \Lambda_{k-m(t)}$ is a step function.

The Λ_i can (at least in theory) be estimated from the volatilities used to value caplets in Black's model (i.e., from the spot volatilities in Figure 28.3).⁷ Suppose that σ_k is the Black volatility for the caplet that corresponds to the period between times t_k and t_{k+1} . Equating variances, we must have

$$\sigma_k^2 t_k = \sum_{i=1}^k \Lambda_{k-i}^2 \delta_{i-1} \quad (31.11)$$

This equation can be used to obtain the Λ 's iteratively.

Example 31.1

Assume that the δ_i are all equal and the Black caplet spot volatilities for the first three caplets are 24%, 22%, and 20%. This means that $\Lambda_0 = 24\%$. Since

$$\Lambda_0^2 + \Lambda_1^2 = 2 \times 0.22^2$$

Λ_1 is 19.80%. Also, since

$$\Lambda_0^2 + \Lambda_1^2 + \Lambda_2^2 = 3 \times 0.20^2$$

Λ_2 is 15.23%.

Example 31.2

Consider the data in Table 31.1 on caplet volatilities σ_k . These exhibit the hump discussed in Section 28.2. The Λ 's are shown in the second row. Notice that the hump in the Λ 's is more pronounced than the hump in the σ 's.

⁶ Since the v 's and ζ 's have opposite signs, the bond price volatility becomes larger (in absolute terms) as the time to maturity increases. This is as expected.

⁷ In practice the Λ 's are determined using a least-squares calibration that we will discuss later.

Table 31.1 Volatility data; accrual period = 1 year.

Year, k :	1	2	3	4	5	6	7	8	9	10
σ_k (%):	15.50	18.25	17.91	17.74	17.27	16.79	16.30	16.01	15.76	15.54
Λ_{k-1} (%):	15.50	20.64	17.21	17.22	15.25	14.15	12.98	13.81	13.60	13.40

Implementation of the Model

The LIBOR market model can be implemented using Monte Carlo simulation. Expressed in terms of the Λ_i 's, equation (31.10) is

$$\frac{dF_k(t)}{F_k(t)} = \sum_{i=m(t)}^k \frac{\delta_i F_i(t) \Lambda_{i-m(t)} \Lambda_{k-m(t)}}{1 + \delta_i F_i(t)} dt + \Lambda_{k-m(t)} dz \quad (31.12)$$

so that from Itô's lemma

$$d \ln F_k(t) = \left[\sum_{i=m(t)}^k \frac{\delta_i F_i(t) \Lambda_{i-m(t)} \Lambda_{k-m(t)}}{1 + \delta_i F_i(t)} - \frac{(\Lambda_{k-m(t)})^2}{2} \right] dt + \Lambda_{k-m(t)} dz \quad (31.13)$$

If, as an approximation, we assume in the calculation of the drift of $\ln F_k(t)$ that $F_i(t) = F_i(t_j)$ for $t_j < t < t_{j+1}$, then

$$F_k(t_{j+1}) = F_k(t_j) \exp \left[\left(\sum_{i=j+1}^k \frac{\delta_i F_i(t_j) \Lambda_{i-j-1} \Lambda_{k-j-1}}{1 + \delta_i F_i(t_j)} - \frac{\Lambda_{k-j-1}^2}{2} \right) \delta_j + \Lambda_{k-j-1} \epsilon \sqrt{\delta_j} \right] \quad (31.14)$$

where ϵ is a random sample from a normal distribution with mean equal to zero and standard deviation equal to one. In the Monte Carlo simulation, this equation is used to calculate forward rates at time t_1 from those at time zero; it is then used to calculate forward rates at time t_2 from those at time t_1 ; and so on.

Extension to Several Factors

The LIBOR market model can be extended to incorporate several independent factors. Suppose that there are p factors and $\zeta_{k,q}$ is the component of the volatility of $F_k(t)$ attributable to the q th factor. Equation (31.10) becomes (see Problem 31.11)

$$\frac{dF_k(t)}{F_k(t)} = \sum_{i=m(t)}^k \frac{\delta_i F_i(t) \sum_{q=1}^p \zeta_{i,q}(t) \zeta_{k,q}(t)}{1 + \delta_i F_i(t)} dt + \sum_{q=1}^p \zeta_{k,q}(t) dz_q \quad (31.15)$$

Define $\lambda_{i,q}$ as the q th component of the volatility when there are i accrual periods between the next reset date and the maturity of the forward contract. Equation (31.14)

then becomes

$$F_k(t_{j+1}) = F_k(t_j) \times \exp \left[\left(\sum_{i=j+1}^k \frac{\delta_i F_i(t_j) \sum_{q=1}^p \lambda_{i-j-1,q} \lambda_{k-j-1,q}}{1 + \delta_i F_i(t_j)} - \frac{\sum_{q=1}^p \lambda_{k-j-1,q}^2}{2} \right) \delta_j + \sum_{q=1}^p \lambda_{k-j-1,q} \epsilon_q \sqrt{\delta_j} \right] \quad (31.16)$$

where the ϵ_q are random samples from a normal distribution with mean equal to zero and standard deviation equal to one.

The approximation that the drift of a forward rate remains constant within each accrual period allows us to jump from one reset date to the next in the simulation. This is convenient because as already mentioned the rolling forward risk-neutral world allows us to discount from one reset date to the next. Suppose that we wish to simulate a zero curve for N accrual periods. On each trial we start with the forward rates at time zero. These are $F_0(0), F_1(0), \dots, F_{N-1}(0)$ and are calculated from the initial zero curve. Equation (31.16) is used to calculate $F_1(t_1), F_2(t_1), \dots, F_{N-1}(t_1)$. Equation (31.16) is then used again to calculate $F_2(t_2), F_3(t_2), \dots, F_{N-1}(t_2)$, and so on, until $F_{N-1}(t_{N-1})$ is obtained. Note that as we move through time the zero curve gets shorter and shorter. For example, suppose each accrual period is 3 months and $N = 40$. We start with a 10-year zero curve. At the 6-year point (at time t_{24}), the simulation gives us information on a 4-year zero curve.

The drift approximation that we have used (i.e., $F_i(t) = F_i(t_j)$ for $t_j < t < t_{j+1}$) can be tested by valuing caplets using equation (31.16) and comparing the prices to those given by Black's model. The value of $F_k(t_k)$ is the realized rate for the time period between t_k and t_{k+1} and enables the caplet payoff at time t_{k+1} to be calculated. This payoff is discounted back to time zero, one accrual period at a time. The caplet value is the average of the discounted payoffs. The results of this type of analysis show that the cap values from Monte Carlo simulation are not significantly different from those given by Black's model. This is true even when the accrual periods are 1 year in length and a

Table 31.2 Valuation of ratchet caplets.

Caplet start time (years)	One factor	Two factors	Three factors
1	0.196	0.194	0.195
2	0.207	0.207	0.209
3	0.201	0.205	0.210
4	0.194	0.198	0.205
5	0.187	0.193	0.201
6	0.180	0.189	0.193
7	0.172	0.180	0.188
8	0.167	0.174	0.182
9	0.160	0.168	0.175
10	0.153	0.162	0.169

Table 31.3 Valuation of sticky caplets.

Caplet start time (years)	One factor	Two factors	Three factors
1	0.196	0.194	0.195
2	0.336	0.334	0.336
3	0.412	0.413	0.418
4	0.458	0.462	0.472
5	0.484	0.492	0.506
6	0.498	0.512	0.524
7	0.502	0.520	0.533
8	0.501	0.523	0.537
9	0.497	0.523	0.537
10	0.488	0.519	0.534

very large number of trials is used.⁸ This suggests that the drift approximation is innocuous in most situations.

Ratchet Caps, Sticky Caps, and Flexi Caps

The LIBOR market model can be used to value some types of nonstandard caps. Consider ratchet caps and sticky caps. These incorporate rules for determining how the cap rate for each caplet is set. In a *ratchet cap* it equals the LIBOR rate at the previous reset date plus a spread. In a *sticky cap* it equals the previous capped rate plus a spread. Suppose that the cap rate at time t_j is K_j , the LIBOR rate at time t_j is R_j , and the spread is s . In a ratchet cap, $K_{j+1} = R_j + s$. In a sticky cap, $K_{j+1} = \min(R_j, K_j) + s$.

Tables 31.2 and 31.3 provide valuations of a ratchet cap and sticky cap using the LIBOR market model with one, two, and three factors. The principal is \$100. The term structure is assumed to be flat at 5% per annum and the caplet volatilities are as in Table 31.1. The interest rate is reset annually. The spread is 25 basis points. Tables 31.4 and 31.5 show how the volatility was split into components when two- and three-factor

Table 31.4 Volatility components in two-factor model.

Year, k :	1	2	3	4	5	6	7	8	9	10
$\lambda_{k-1,1}$ (%):	14.10	19.52	16.78	17.11	15.25	14.06	12.65	13.06	12.36	11.63
$\lambda_{k-1,2}$ (%):	-6.45	-6.70	-3.84	-1.96	0.00	1.61	2.89	4.48	5.65	6.65
Total volatility (%):	15.50	20.64	17.21	17.22	15.25	14.15	12.98	13.81	13.60	13.40

⁸ See J. C. Hull and A. White, "Forward Rate Volatilities, Swap Rate Volatilities, and the Implementation of the LIBOR Market Model," *Journal of Fixed Income*, 10, 2 (September 2000): 46-62. The only exception is when the cap volatilities are very high.

Table 31.5 Volatility components in a three-factor model.

Year, k :	1	2	3	4	5	6	7	8	9	10
$\lambda_{k-1,1}$ (%):	13.65	19.28	16.72	16.98	14.85	13.95	12.61	12.90	11.97	10.97
$\lambda_{k-1,2}$ (%):	-6.62	-7.02	-4.06	-2.06	0.00	1.69	3.06	4.70	5.81	6.66
$\lambda_{k-1,3}$ (%):	3.19	2.25	0.00	-1.98	-3.47	-1.63	0.00	1.51	2.80	3.84
Total volatility (%):	15.50	20.64	17.21	17.22	15.25	14.15	12.98	13.81	13.60	13.40

models were used. The results are based on 100,000 Monte Carlo simulations incorporating the antithetic variable technique described in Section 20.7. The standard error of each price is about 0.001.

A third type of nonstandard cap is a *flexi cap*. This is like a regular cap except that there is a limit on the total number of caplets that can be exercised. Consider an annual-pay flexi cap when the principal is \$100, the term structure is flat at 5%, and the cap volatilities are as in Tables 31.1, 31.4, and 31.5. Suppose that all in-the-money caplets are exercised up to a maximum of five. With one, two, and three factors, the LIBOR market model gives the price of the instrument as 3.43, 3.58, and 3.61, respectively (see Problem 31.15 for other types of flexi caps).

The pricing of a plain vanilla cap depends only on the total volatility and is independent of the number of factors. This is because the price of a plain vanilla caplet depends on the behavior of only one forward rate. The prices of caplets in the nonstandard instruments we have looked at are different in that they depend on the joint probability distribution of several different forward rates. As a result they do depend on the number of factors.

Valuing European Swap Options

As shown by Hull and White, there is an analytic approximation for valuing European swap options in the LIBOR market model.⁹ Let T_0 be the maturity of the swap option and assume that the payment dates for the swap are T_1, T_2, \dots, T_N . Define $\tau_i = T_{i+1} - T_i$. From equation (27.23), the swap rate at time t is given by

$$s(t) = \frac{P(t, T_0) - P(t, T_N)}{\sum_{i=0}^{N-1} \tau_i P(t, T_{i+1})}$$

It is also true that

$$\frac{P(t, T_i)}{P(t, T_0)} = \prod_{j=0}^{i-1} \frac{1}{1 + \tau_j G_j(t)}$$

⁹ See J. C. Hull and A. White, "Forward Rate Volatilities, Swap Rate Volatilities, and the Implementation of the LIBOR Market Model," *Journal of Fixed Income*, 10, 2 (September 2000): 46–62. Other analytic approximations have been suggested by A. Brace, D. Gatarek, and M. Musiela "The Market Model of Interest Rate Dynamics," *Mathematical Finance*, 7, 2 (1997): 127–55 and L. Andersen and J. Andreasen, "Volatility Skews and Extensions of the LIBOR Market Model," *Applied Mathematical Finance*, 7, 1 (March 2000), 1–32.

for $1 \leq i \leq N$, where $G_j(t)$ is the forward rate at time t for the period between T_j and T_{j+1} . These two equations together define a relationship between $s(t)$ and the $G_j(t)$. Applying Itô's lemma (see Problem 31.12), the variance $V(t)$ of the swap rate $s(t)$ is given by

$$V(t) = \sum_{q=1}^p \left[\sum_{k=0}^{N-1} \frac{\tau_k \beta_{k,q}(t) G_k(t) \gamma_k(t)}{1 + \tau_k G_k(t)} \right]^2 \quad (31.17)$$

where

$$\gamma_k(t) = \frac{\prod_{j=0}^{N-1} [1 + \tau_j G_j(t)]}{\prod_{j=0}^{N-1} [1 + \tau_j G_j(t)] - 1} - \frac{\sum_{i=0}^{k-1} \tau_i \prod_{j=i+1}^{N-1} [1 + \tau_j G_j(t)]}{\sum_{i=0}^{N-1} \tau_i \prod_{j=i+1}^N [1 + \tau_j G_j(t)]}$$

and $\beta_{j,q}(t)$ is the q th component of the volatility of $G_j(t)$. We approximate $V(t)$ by setting $G_j(t) = G_j(0)$ for all j and t . The swap volatility that is substituted into the standard market model for valuing a swaption is then

$$\sqrt{\frac{1}{T_0} \int_{t=0}^{T_0} V(t) dt}$$

or

$$\sqrt{\frac{1}{T_0} \int_{t=0}^{T_0} \sum_{q=1}^p \left[\sum_{k=0}^{N-1} \frac{\tau_k \beta_{k,q}(t) G_k(0) \gamma_k(0)}{1 + \tau_k G_k(0)} \right]^2 dt} \quad (31.18)$$

In the situation where the length of the accrual period for the swap underlying the swaption is the same as the length of the accrual period for a cap, $\beta_{k,q}(t)$ is the q th component of volatility of a cap forward rate when the time to maturity is $T_k - t$. This can be looked up in a table such as Table 31.5

The accrual periods for the swaps underlying broker quotes for European swap options do not always match the accrual periods for the caps and floors underlying broker quotes. For example, in the United States, the benchmark caps and floors have quarterly resets, while the swaps underlying the benchmark European swap options have semiannual resets. Fortunately, the valuation result for European swap options can be extended to the situation where each swap accrual period includes M subperiods that could be accrual periods in a typical cap. Define $\tau_{j,m}$ as the length of the m th subperiod in the j th accrual period so that

$$\tau_j = \sum_{m=1}^M \tau_{j,m}$$

Define $G_{j,m}(t)$ as the forward rate observed at time t for the $\tau_{j,m}$ accrual period. Because

$$1 + \tau_j G_j(t) = \prod_{m=1}^M [1 + \tau_{j,m} G_{j,m}(t)]$$

the analysis leading to equation (31.18) can be modified so that the volatility of $s(t)$ is obtained in terms of the volatilities of the $G_{j,m}(t)$ rather than the volatilities of the $G_j(t)$.

The swap volatility to be substituted into the standard market model for valuing a swap option proves to be (see Problem 31.13)

$$\sqrt{\frac{1}{T_0} \int_{t=0}^{T_0} \sum_{q=1}^p \left[\sum_{k=n}^{N-1} \sum_{m=1}^M \frac{\tau_{k,m} \beta_{k,m,q}(t) G_{k,m}(0) \gamma_k(0)}{1 + \tau_{k,m} G_{k,m}(0)} \right]^2 dt} \quad (31.19)$$

Here $\beta_{j,m,q}(t)$ is the q th component of the volatility of $G_{j,m}(t)$. It is the q th component of the volatility of a cap forward rate when the time to maturity is from t to the beginning of the m th subperiod in the (T_j, T_{j+1}) swap accrual period.

The expressions in equations (31.18) and (31.19) for the swap volatility do involve the approximations that $G_j(t) = G_j(0)$ and $G_{j,m}(t) = G_{j,m}(0)$. Hull and White compared the prices of European swap options calculated using equations (31.18) and (31.19) with the prices calculated from a Monte Carlo simulation and found the two to be very close. Once the LIBOR market model has been calibrated, equations (31.18) and (31.19) therefore provide a quick way of valuing European swap options. Analysts can determine whether European swap options are overpriced or underpriced relative to caps. As we will see shortly, they can also use the results to calibrate the model to the market prices of swap options.

Calibrating the Model

The variable Λ_j is the volatility at time t of the forward rate F_j for the period between t_k and t_{k+1} when there are j whole accrual periods between t and t_k . To calibrate the LIBOR market model, it is necessary to determine the Λ_j and how they are split into $\lambda_{j,q}$. The Λ 's are usually determined from current market data, whereas the split into λ 's is determined from historical data.

Consider first the determination of the λ 's from the Λ 's. A principal components analysis (see Section 21.9) on forward rate data can be used. The model is

$$\Delta F_j = \sum_{q=1}^M \alpha_{j,q} x_q$$

where M is the total number of factors (which equals the number of different forward rates), ΔF_j is the change in the j th forward rate F_j , $\alpha_{j,q}$ is the factor loading for the j th forward rate and the q th factor, x_q is the factor score for the q th factor. Define s_q as the standard deviation of the q th factor score. If the number of factors used in the LIBOR market model, p , is equal to the total number of factors, M , it is correct to set

$$\lambda_{j,q} = \alpha_{j,q} s_q$$

for $1 \leq j, q \leq M$. When $p < M$, the $\lambda_{j,q}$ must be scaled so that

$$\Lambda_j = \sqrt{\sum_{q=1}^p \lambda_{j,q}^2}$$

This involves setting

$$\lambda_{j,q} = \frac{\Lambda_j s_q \alpha_{j,q}}{\sqrt{\sum_{q=1}^p s_q^2 \alpha_{j,q}^2}} \quad (31.20)$$

Consider next the estimation of the Λ 's. Equation (31.11) provides one way that they can be theoretically determined so that they are consistent with caplet prices. In practice, this is not usually used because it often leads to wild swings in the Λ 's and sometimes there is no set of Λ 's exactly consistent with cap quotes. A commonly used calibration procedure is similar to that described for one-factor models in Section 30.8. Suppose that U_i is the market price of the i th calibrating instrument (typically a cap or European swaption) and V_i is the model price. The Λ 's are chosen to minimize

$$\sum_i (U_i - V_i)^2 + P$$

where P is a penalty function chosen to ensure that the Λ 's are "well behaved." Similarly to Section 30.8, P might have the form

$$P = \sum_i w_{1,i} (\Lambda_{i+1} - \Lambda_i)^2 + \sum_i w_{2,i} (\Lambda_{i+1} + \Lambda_{i-1} - 2\Lambda_i)^2$$

When the calibrating instrument is a European swaption, formulas (31.18) and (31.19) make the minimization feasible using the Levenberg–Marquardt procedure. Equation (31.20) is used to determine the λ 's from the Λ 's.

Volatility Skews

Brokers provide quotes on caps that are not at the money as well as on caps that are at the money. In some markets a volatility skew is observed, that is, the quoted (Black) volatility for a cap or a floor is a declining function of the strike price. This can be handled using the CEV model. (See Section 26.1 for the application of the CEV model to equities.) The model is

$$dF_i(t) = \cdots + \sum_{q=1}^p \zeta_{i,q}(t) F_i(t)^\alpha dz_q \quad (31.21)$$

where α is a constant ($0 < \alpha < 1$). It turns out that this model can be handled very similarly to the lognormal model. Caps and floors can be valued analytically using the cumulative noncentral χ^2 distribution. There are similar analytic approximations to those given above for the prices of European swap options.¹⁰

Bermudan Swap Options

A popular interest rate derivative is a Bermudan swap option. This is a swap option that can be exercised on some or all of the payment dates of the underlying swap. Bermudan swap options are difficult to value using the LIBOR market model because the LIBOR market model relies on Monte Carlo simulation and it is difficult to evaluate early exercise decisions when Monte Carlo simulation is used. Fortunately, the procedures described in Section 26.8 can be used. Longstaff and Schwartz apply the least-squares approach when there are a large number of factors. The value of not

¹⁰ For details, see L. Andersen and J. Andreasen, "Volatility Skews and Extensions of the LIBOR Market Model," *Applied Mathematical Finance*, 7, 1 (2000): 1–32; J.C. Hull and A. White, "Forward Rate Volatilities, Swap Rate Volatilities, and the Implementation of the LIBOR Market Model," *Journal of Fixed Income*, 10, 2 (September 2000): 46–62.

exercising on a particular payment date is assumed to be a polynomial function of the values of the factors.¹¹ Andersen shows that the optimal early exercise boundary approach can be used. He experiments with a number of ways of parameterizing the early exercise boundary and finds that good results are obtained when the early exercise decision is assumed to depend only on the intrinsic value of the option.¹² Most traders value Bermudan options using one of the one-factor no-arbitrage models discussed in Chapter 30. However, the accuracy of one-factor models for pricing Bermudan swap options has been a controversial issue.¹³

31.3 AGENCY MORTGAGE-BACKED SECURITIES

One application of the models presented in this chapter is to the agency mortgage-backed security (agency MBS) market in the United States.

An agency MBS is similar to the ABS considered in Chapter 8 except that payments are guaranteed by a government-related agency such as the Government National Mortgage Association (GNMA) or the Federal National Mortgage Association (FNMA) so that investors are protected against defaults. This makes an agency MBS sound like a regular fixed-income security issued by the government. In fact, there is a critical difference between an agency MBS and a regular fixed-income investment. This difference is that the mortgages in an agency MBS pool have prepayment privileges. These prepayment privileges can be quite valuable to the householder. In the United States, mortgages typically last for 30 years and can be prepaid at any time. This means that the householder has a 30-year American-style option to put the mortgage back to the lender at its face value.

Prepayments on mortgages occur for a variety of reasons. Sometimes interest rates fall and the owner of the house decides to refinance at a lower rate. On other occasions, a mortgage is prepaid simply because the house is being sold. A critical element in valuing an agency MBS is the determination of what is known as the *prepayment function*. This is a function describing expected prepayments on the underlying pool of mortgages at a time t in terms of the yield curve at time t and other relevant variables.

A prepayment function is very unreliable as a predictor of actual prepayment experience for an individual mortgage. When many similar mortgage loans are combined in the same pool, there is a “law of large numbers” effect at work and prepayments can be predicted more accurately from an analysis of historical data. As mentioned, prepayments are not always motivated by pure interest rate considerations. Nevertheless, there is a tendency for prepayments to be more likely when interest rates are low than when they are high. This means that investors require a higher rate of interest on an agency MBS than on other fixed-income securities to compensate for the prepayment options they have written.

¹¹ See F. A. Longstaff and E. S. Schwartz, “Valuing American Options by Simulation: A Simple Least Squares Approach,” *Review of Financial Studies*, 14, 1 (2001): 113–47.

¹² L. Andersen, “A Simple Approach to the Pricing of Bermudan Swaptions in the Multifactor LIBOR Market Model,” *Journal of Computational Finance*, 3, 2 (Winter 2000): 5–32.

¹³ For opposing viewpoints, see “Factor Dependence of Bermudan Swaptions: Fact or Fiction,” by L. Andersen and J. Andreasen, and “Throwing Away a Billion Dollars: The Cost of Suboptimal Exercise Strategies in the Swaption Market,” by F. A. Longstaff, P. Santa-Clara, and E. S. Schwartz. Both articles are in *Journal of Financial Economics*, 62, 1 (October 2001).

Collateralized Mortgage Obligations

The simplest type of agency MBS is referred to as a *pass-through*. All investors receive the same return and bear the same prepayment risk. Not all mortgage-backed securities work in this way. In a *collateralized mortgage obligation* (CMO) the investors are divided into a number of classes and rules are developed for determining how principal repayments are channeled to different classes. A CMO creates classes of securities that bear different amounts of prepayment risk in the same way that the ABS considered in Chapter 8 creates classes of securities bearing different amounts of credit risk.

As an example of a CMO, consider an agency MBS where investors are divided into three classes: class A, class B, and class C. All the principal repayments (both those that are scheduled and those that are prepayments) are channeled to class A investors until investors in this class have been completely paid off. Principal repayments are then channeled to class B investors until these investors have been completely paid off. Finally, principal repayments are channeled to class C investors. In this situation, class A investors bear the most prepayment risk. The class A securities can be expected to last for a shorter time than the class B securities, and these, in turn, can be expected to last less long than the class C securities.

The objective of this type of structure is to create classes of securities that are more attractive to institutional investors than those created by a simpler pass-through MBS. The prepayment risks assumed by the different classes depend on the par value in each class. For example, class C bears very little prepayment risk if the par values in classes A, B, and C are 400, 300, and 100, respectively. Class C bears rather more prepayment risk in the situation where the par values in the classes are 100, 200, and 500.

The creators of mortgage-backed securities have created many more exotic structures than the one we have just described. Business Snapshot 31.1 gives an example.

Valuing Agency Mortgage-Backed Securities

Agency MBSs are usually valued using Monte Carlo simulation. Either the HJM or LIBOR market models can be used to simulate the behavior of interest rates month by month throughout the life of an agency MBS. Consider what happens on one simulation trial. Each month, expected prepayments are calculated from the current yield curve and the history of yield curve movements. These prepayments determine the expected cash flows to the holder of the agency MBS and the cash flows are discounted to time zero to obtain a sample value for the agency MBS. An estimate of the value of the agency MBS is the average of the sample values over many simulation trials.

Option-Adjusted Spread

In addition to calculating theoretical prices for mortgage-backed securities and other bonds with embedded options, traders also like to compute what is known as the *option-adjusted spread* (OAS). This is a measure of the spread over the yields on government Treasury bonds provided by the instrument when all options have been taken into account.

An input to any term structure model is the initial zero-coupon yield curve. Usually this is the LIBOR zero curve. However, to calculate an OAS for an instrument, it is first priced using the zero-coupon government Treasury curve. The price of the instrument given by the model is compared to the price in the market. A series of iterations is then

Business Snapshot 31.1 IOs and POs

In what is known as a *stripped MBS*, principal payments are separated from interest payments. All principal payments are channeled to one class of security, known as a *principal only* (PO). All interest payments are channeled to another class of security known as an *interest only* (IO). Both IOs and POs are risky investments. As prepayment rates increase, a PO becomes more valuable and an IO becomes less valuable. As prepayment rates decrease, the reverse happens. In a PO, a fixed amount of principal is returned to the investor, but the timing is uncertain. A high rate of prepayments on the underlying pool leads to the principal being received early (which is, of course, good news for the holder of the PO). A low rate of prepayments on the underlying pool delays the return of the principal and reduces the yield provided by the PO. In the case of an IO, the total of the cash flows received by the investor is uncertain. The higher the rate of prepayments, the lower the total cash flows received by the investor, and vice versa.

used to determine the parallel shift to the input Treasury curve that causes the model price to be equal to the market price. This parallel shift is the OAS.

To illustrate the nature of the calculations, suppose that the market price is \$102.00 and that the price calculated using the Treasury curve is \$103.27. As a first trial we might choose to try a 60-basis-point parallel shift to the Treasury zero curve. Suppose that this gives a price of \$101.20 for the instrument. This is less than the market price of \$102.00 and means that a parallel shift somewhere between 0 and 60 basis points will lead to the model price being equal to the market price. We could use linear interpolation to calculate

$$60 \times \frac{103.27 - 102.00}{103.27 - 101.20} = 36.81$$

or 36.81 basis points as the next trial shift. Suppose that this gives a price of \$101.95. This indicates that the OAS is slightly less than 36.81 basis points. Linear interpolation suggests that the next trial shift be

$$36.81 \times \frac{103.27 - 102.00}{103.27 - 101.95} = 35.41$$

or 35.41 basis points; and so on.

SUMMARY

The HJM and LMM models provide approaches to valuing interest rate derivatives that give the user complete freedom in choosing the volatility term structure. The LMM model has two key advantages over the HJM model. First, it is developed in terms of the forward rates that determine the pricing of caps, rather than in terms of instantaneous forward rates. Second, it is relatively easy to calibrate to the price of caps or European swap options. The HJM and LMM models both have the disadvantage that they cannot be represented as recombining trees. In practice, this means that they must be implemented using Monte Carlo simulation.

The agency mortgage-backed security market in the United States has given birth to many exotic interest rate derivatives: CMOs, IOs, POs, and so on. These instruments provide cash flows to the holder that depend on the prepayments on a pool of mortgages. These prepayments depend on, among other things, the level of interest rates. Because they are heavily path dependent, agency mortgage-backed securities usually have to be valued using Monte Carlo simulation. These are, therefore, ideal candidates for applications of the HJM and LMM models.

FURTHER READING

- Andersen, L., “A Simple Approach to the Pricing of Bermudan Swaption in the Multi-Factor LIBOR Market Model,” *The Journal of Computational Finance*, 3, 2 (2000): 5–32.
- Andersen, L., and J. Andreasen, “Volatility Skews and Extensions of the LIBOR Market Model,” *Applied Mathematical Finance*, 7, 1 (March 2000): 1–32.
- Andersen, L., and V. Piterbarg, *Interest Rate Modeling*, Vols. I–III. New York: Atlantic Financial Press, 2010.
- Brace A., D. Gatarek, and M. Musiela “The Market Model of Interest Rate Dynamics,” *Mathematical Finance*, 7, 2 (1997): 127–55.
- Duffie, D. and R. Kan, “A Yield-Factor Model of Interest Rates,” *Mathematical Finance* 6, 4 (1996), 379–406.
- Heath, D., R. Jarrow, and A. Morton, “Bond Pricing and the Term Structure of Interest Rates: A Discrete Time Approximation,” *Journal of Financial and Quantitative Analysis*, 25, 4 (December 1990): 419–40.
- Heath, D., R. Jarrow, and A. Morton, “Bond Pricing and the Term Structure of the Interest Rates: A New Methodology,” *Econometrica*, 60, 1 (1992): 77–105.
- Hull, J., and A. White, “Forward Rate Volatilities, Swap Rate Volatilities, and the Implementation of the LIBOR Market Model,” *Journal of Fixed Income*, 10, 2 (September 2000): 46–62.
- Inui, K., and M. Kijima, “A Markovian Framework in Multifactor Heath, Jarrow, and Morton Models,” *Journal of Financial and Quantitative Analysis*, 33, 3 (September 1998): 423–40.
- Jamshidian, F., “LIBOR and Swap Market Models and Measures,” *Finance and Stochastics*, 1 (1997): 293–330.
- Jarrow, R. A., *Modeling Fixed Income Securities and Interest Rate Options*, 2nd edn. New York: McGraw-Hill, 2002.
- Jarrow, R. A., and S. M. Turnbull, “Delta, Gamma, and Bucket Hedging of Interest Rate Derivatives,” *Applied Mathematical Finance*, 1 (1994): 21–48.
- Miltersen, K., K. Sandmann, and D. Sondermann, “Closed Form Solutions for Term Structure Derivatives with Lognormal Interest Rates,” *Journal of Finance*, 52, 1 (March 1997): 409–30.
- Rebonato, R., *Modern Pricing of Interest Rate Derivatives: The LIBOR Market Model and Beyond*. Princeton University Press, 2002.

Practice Questions (Answers in Solutions Manual)

- 31.1. Explain the difference between a Markov and a non-Markov model of the short rate.
- 31.2. Prove the relationship between the drift and volatility of the forward rate for the multifactor version of HJM in equation (31.6).

- 31.3. “When the forward rate volatility $s(t, T)$ in HJM is constant, the Ho–Lee model results.” Verify that this is true by showing that HJM gives a process for bond prices that is consistent with the Ho–Lee model in Chapter 30.
- 31.4. “When the forward rate volatility, $s(t, T)$, in HJM is $\sigma e^{-a(T-t)}$, the Hull–White model results.” Verify that this is true by showing that HJM gives a process for bond prices that is consistent with the Hull–White model in Chapter 30.
- 31.5. What is the advantage of LMM over HJM?
- 31.6. Provide an intuitive explanation of why a ratchet cap increases in value as the number of factors increase.
- 31.7. Show that equation (31.10) reduces to (31.4) as the δ_i tend to zero.
- 31.8. Explain why a sticky cap is more expensive than a similar ratchet cap.
- 31.9. Explain why IOs and POs have opposite sensitivities to the rate of prepayments.
- 31.10. “An option adjusted spread is analogous to the yield on a bond.” Explain this statement.
- 31.11. Prove equation (31.15).
- 31.12. Prove the formula for the variance $V(T)$ of the swap rate in equation (31.17).
- 31.13. Prove equation (31.19).

Further Questions

- 31.14. In an annual-pay cap, the Black volatilities for caplets with maturities 1, 2, 3, and 5 years are 18%, 20%, 22%, and 20%, respectively. Estimate the volatility of a 1-year forward rate in the LIBOR Market Model when the time to maturity is (a) 0 to 1 year, (b) 1 to 2 years, (c) 2 to 3 years, and (d) 3 to 5 years. Assume that the zero curve is flat at 5% per annum (annually compounded). Use DerivaGem to estimate flat volatilities for 2-, 3-, 4-, 5-, and 6-year caps.
- 31.15. In the flexi cap considered in Section 31.2 the holder is obligated to exercise the first N in-the-money caplets. After that no further caplets can be exercised. (In the example, $N = 5$.) Two other ways that flexi caps are sometimes defined are:
 - (a) The holder can choose whether any caplet is exercised, but there is a limit of N on the total number of caplets that can be exercised.
 - (b) Once the holder chooses to exercise a caplet all subsequent in-the-money caplets must be exercised up to a maximum of N .Discuss the problems in valuing these types of flexi caps. Of the three types of flexi caps, which would you expect to be most expensive? Which would you expect to be least expensive?



CHAPTER 32

Swaps Revisited

Swaps have been central to the success of over-the-counter derivatives markets. They have proved to be very flexible instruments for managing risk. Based on the range of different contracts that now trade and the total volume of business transacted each year, swaps are arguably one of the most successful innovations in financial markets ever.

Chapter 7 discussed how plain vanilla interest rate swaps can be valued. The standard approach can be summarized as: “Assume forward rates will be realized.” The steps are as follows:

1. Calculate the swap’s net cash flows on the assumption that LIBOR rates in the future equal the forward rates calculated from today’s LIBOR/swap zero curve.
2. Set the value of the swap equal to the present value of the net cash flows using the LIBOR/swap zero curve for discounting.

This chapter describes a number of nonstandard swaps. Some can be valued using the “assume forward rates will be realized” approach; some require the application of the convexity, timing, and quanto adjustments we encountered in Chapters 29; some contain embedded options that must be valued using the procedures described in Chapters 28, 30, and 31.

32.1 VARIATIONS ON THE VANILLA DEAL

Many interest rate swaps involve relatively minor variations to the plain vanilla structure discussed in Chapter 7. In some swaps the notional principal changes with time in a predetermined way. Swaps where the notional principal is an increasing function of time are known as *step-up swaps*. Swaps where the notional principal is a decreasing function of time are known as *amortizing swaps*. Step-up swaps could be useful for a construction company that intends to borrow increasing amounts of money at floating rates to finance a particular project and wants to swap to fixed-rate funding. An amortizing swap could be used by a company that has fixed-rate borrowings with a certain prepayment schedule and wants to swap to borrowings at a floating rate.

Business Snapshot 32.1 Hypothetical Confirmation for Nonstandard Swap

Trade date:	5-January, 2010
Effective date:	11-January, 2010
Business day convention (all dates):	Following business day
Holiday calendar:	US
Termination date:	11-January, 2015
<i>Fixed amounts</i>	
Fixed-rate payer:	Microsoft
Fixed-rate notional principal:	USD 100 million
Fixed rate:	6% per annum
Fixed-rate day count convention:	Actual/365
Fixed-rate payment dates	Each 11-July and 11-January commencing 11-July, 2010, up to and including 11-January, 2015
<i>Floating amounts</i>	
Floating-rate payer	Goldman Sachs
Floating-rate notional principal	USD 120 million
Floating rate	USD 1-month LIBOR
Floating-rate day count convention	Actual/360
Floating-rate payment dates	11-July, 2010, and the 11th of each month thereafter up to and including 11-January, 2015

The principal can be different on the two sides of a swap. Also the frequency of payments can be different. Business Snapshot 32.1 illustrates this by showing a hypothetical swap between Microsoft and Goldman Sachs where the notional principal is \$120 million on the floating side and \$100 million on fixed side. Payments are made every month on the floating side and every 6 months on the fixed side. These type of variations to the basic plain vanilla structure do not affect the valuation methodology. The “assume forward rates are realized” approach can still be used.

The floating reference rate for a swap is not always LIBOR. In some swaps for instance, it is the commercial paper (CP) rate. A *basis swap* involves exchanging cash flows calculated using one floating reference rate for cash flows calculated using another floating reference rate. An example would be a swap where the 3-month CP rate plus 10 basis points is exchanged for 3-month LIBOR with both being applied to a principal of \$100 million. A basis swap could be used for risk management by a financial institution whose assets and liabilities are dependent on different floating reference rates.

Swaps where the floating reference rate is not LIBOR can be valued using the “assume forward rates are realized” approach. A zero curve other than LIBOR is necessary to calculate future cash flows on the assumption that forward rates are realized. The cash flows are discounted at LIBOR.

Business Snapshot 32.2 Hypothetical Confirmation for Compounding Swap

Trade date:	5-January, 2010
Effective date:	11-January, 2010
Holiday calendar:	US
Business day convention (all dates):	Following business day
Termination date:	11-January, 2015
<i>Fixed amounts</i>	
Fixed-rate payer:	Microsoft
Fixed-rate notional principal:	USD 100 million
Fixed rate:	6% per annum
Fixed-rate day count convention:	Actual/365
Fixed-rate payment date:	11-January, 2015
Fixed-rate compounding:	Applicable at 6.3%
Fixed-rate compounding dates	Each 11-July and 11-January commencing 11-July, 2010, up to and including 11-July, 2014
<i>Floating amounts</i>	
Floating-rate payer:	Goldman Sachs
Floating-rate notional principal:	USD 100 million
Floating rate:	USD 6-month LIBOR plus 20 basis points
Floating-rate day count convention:	Actual/360
Floating-rate payment date:	11-January, 2015
Floating-rate compounding:	Applicable at LIBOR plus 10 basis points
Floating-rate compounding dates:	Each 11-July and 11-January commencing 11-July, 2010, up to and including 11-July, 2014

32.2 COMPOUNDING SWAPS

Another variation on the plain vanilla swap is a *compounding swap*. A hypothetical confirmation for a compounding swap is in Business Snapshot 32.2. In this example there is only one payment date for both the floating-rate payments and the fixed-rate payments. This is at the end of the life of the swap. The floating rate of interest is LIBOR plus 20 basis points. Instead of being paid, the interest is compounded forward until the end of the life of the swap at a rate of LIBOR plus 10 basis points. The fixed rate of interest is 6%. Instead of being paid this interest is compounded forward at a fixed rate of interest of 6.3% until the end of the swap.

The “assume forward rates are realized” approach can be used at least approximately for valuing a compounding swap such as that in Business Snapshot 32.2. It is straightforward to deal with the fixed side of the swap because the payment that will be made at maturity is known with certainty. The “assume forward rates are realized” approach for

the floating side is justifiable because there exist a series of forward rate agreements (FRAs) where the floating-rate cash flows are exchanged for the values they would have if each floating rate equaled the corresponding forward rate.¹

Example 32.1

A compounding swap with annual resets has a life of 3 years. A fixed rate is paid and a floating rate is received. The fixed interest rate is 4% and the floating interest rate is 12-month LIBOR. The fixed side compounds at 3.9% and the floating side compounds at 12-month LIBOR minus 20 basis points. The LIBOR zero curve is flat at 5% with annual compounding and the notional principal is \$100 million.

On the fixed side, interest of \$4 million is earned at the end of the first year. This compounds to $4 \times 1.039 = \$4.156$ million at the end of the second year. A second interest amount of \$4 million is added at the end of the second year bringing the total compounded forward amount to \$8.156 million. This compounds to $8.156 \times 1.039 = \$8.474$ million by the end of the third year when there is the third interest amount of \$4 million. The cash flow at the end of the third year on the fixed side of the swap is therefore \$12.474 million.

On the floating side we assume all future interest rates equal the corresponding forward LIBOR rates. Given the LIBOR zero curve, this means that all future interest rates are assumed to be 5% with annual compounding. The interest calculated at the end of the first year is \$5 million. Compounding this forward at 4.8% (forward LIBOR minus 20 basis points) gives $5 \times 1.048 = \$5.24$ million at the end of the second year. Adding in the interest, the compounded forward amount is \$10.24 million. Compounding forward to the end of the third year, we get $10.24 \times 1.048 = \$10.731$ million. Adding in the final interest gives \$15.731 million.

The swap can be valued by assuming that it leads to an inflow of \$15.731 million and an outflow of \$12.474 million at the end of year 3. The value of the swap is therefore

$$\frac{15.731 - 12.474}{1.05^3} = 2.814$$

or \$2.814 million. (This analysis ignores day count issues and makes the approximation indicated in footnote 1.)

32.3 CURRENCY SWAPS

Currency swaps were introduced in Chapter 7. They enable an interest rate exposure in one currency to be swapped for an interest rate exposure in another currency. Usually two principals are specified, one in each currency. The principals are exchanged at both the beginning and the end of the life of the swap as described in Section 7.8.

Suppose that the currencies involved in a currency swap are US dollars (USD) and British pounds (GBP). In a fixed-for-fixed currency swap, a fixed rate of interest is specified in each currency. The payments on one side are determined by applying the

¹ See Technical Note 18 at www.rotman.utoronto.ca/~hull/TechnicalNotes for the details. The “assume forward rates are realized” approach works exactly if the spread used for compounding, s_c , is zero or if it is applied so that Q at time t compounds to $Q(1 + R\tau)(1 + s_c\tau)$ at time $t + \tau$, where R is LIBOR. If, as is more usual, it compounds to $Q[1 + (R + s_c)\tau]$, then there is a small approximation.

fixed rate of interest in USD to the USD principal; the payments on the other side are determined by applying the fixed rate of interest in GBP to the GBP principal. Section 7.9 discussed the valuation of this type of swap.

Another popular type of currency swap is floating-for-floating. In this, the payments on one side are determined by applying USD LIBOR (possibly with a spread added) to the USD principal; similarly the payments on the other side are determined by applying GBP LIBOR (possibly with a spread added) to the GBP principal. A third type of swap is a cross-currency interest rate swap where a floating rate in one currency is exchanged for a fixed rate in another currency.

Floating-for-floating and cross-currency interest rate swaps can be valued using the “assume forward rates are realized” rule. Future LIBOR rates in each currency are assumed to equal today’s forward rates. This enables the cash flows in the currencies to be determined. The USD cash flows are discounted at the USD LIBOR zero rate. The GBP cash flows are discounted at the GBP LIBOR zero rate. The current exchange rate is then used to translate the two present values to a common currency.

An adjustment to this procedure is sometimes made to reflect the realities of the market. In theory, a new floating-for-floating swap should involve exchanging LIBOR in one currency for LIBOR in another currency (with no spreads added). In practice, macroeconomic effects give rise to spreads. Financial institutions often adjust the discount rates they use to allow for this. As an example, suppose that market conditions are such that USD LIBOR is exchanged for Japanese yen (JPY) LIBOR minus 20 basis points in new floating-for-floating swaps of all maturities. In its valuations a US financial institution would discount USD cash flows at USD LIBOR and it would discount JPY cash flows at JPY LIBOR minus 20 basis points.² It would do this in all swaps that involved both JPY and USD cash flows.

32.4 MORE COMPLEX SWAPS

We now move on to consider some examples of swaps where the simple rule “assume forward rates will be realized” does not work. In each case, it must be assumed that an adjusted forward rate, rather than the actual forward rate, is realized. This section builds on the discussion in Chapter 29.

LIBOR-in-Arrears Swap

A plain vanilla interest rate swap is designed so that the floating rate of interest observed on one payment date is paid on the next payment date. An alternative instrument that is sometimes traded is a *LIBOR-in-arrears swap*. In this, the floating rate paid on a payment date equals the rate observed on the payment date itself.

Suppose that the reset dates in the swap are t_i for $i = 0, 1, \dots, n$, with $\tau_i = t_{i+1} - t_i$. Define R_i as the LIBOR rate for the period between t_i and t_{i+1} , F_i as the forward value of R_i , and σ_i as the volatility of this forward rate. (The value of σ_i is typically implied from caplet prices.) In a LIBOR-in-arrears swap the payment on the floating side at time t_i is based on R_i rather than R_{i-1} . As explained in Section 29.1, it is necessary to

² This adjustment is *ad hoc*, but, if it is not made, traders make an immediate profit or loss every time they trade a new JPY/USD floating-for-floating swap.

make a convexity adjustment to the forward rate when the payment is valued. The valuation should be based on the assumption that the floating rate paid is

$$F_i + \frac{F_i^2 \sigma_i^2 \tau_i t_i}{1 + F_i \tau_i} \quad (32.1)$$

and not F_i .

Example 32.2

In a LIBOR-in-arrears swap, the principal is \$100 million. A fixed rate of 5% is received annually and LIBOR is paid. Payments are exchanged at the ends of years 1, 2, 3, 4, and 5. The yield curve is flat at 5% per annum (measured with annual compounding). All caplet volatilities are 22% per annum.

The forward rate for each floating payment is 5%. If this were a regular swap rather than an in-arrears swap, its value would (ignoring day count conventions, etc.) be exactly zero. Because it is an in-arrears swap, convexity adjustments must be made. In equation (32.1), $F_i = 0.05$, $\sigma_i = 0.22$, and $\tau_i = 1$ for all i . The convexity adjustment changes the rate assumed at time t_i from 0.05 to

$$0.05 + \frac{0.05^2 \times 0.22^2 \times 1 \times t_i}{1 + 0.05 \times 1} = 0.05 + 0.000115t_i$$

The floating rates for the payments at the ends of years 1, 2, 3, 4, and 5 should therefore be assumed to be 5.0115%, 5.0230%, 5.0345%, 5.0460%, and 5.0575%, respectively. The net exchange on the first payment date is equivalent to a cash outflow of 0.0115% of \$100 million or \$11,500. Equivalent net cash flows for other exchanges are calculated similarly. The value of the swap is

$$\begin{aligned} & -\frac{11,500}{1.05} - \frac{23,000}{1.05^2} - \frac{34,500}{1.05^3} - \frac{46,000}{1.05^4} - \frac{57,500}{1.05^5} \\ & \text{or } -\$144,514. \end{aligned}$$

CMS and CMT Swaps

A constant maturity swap (CMS) is an interest rate swap where the floating rate equals the swap rate for a swap with a certain life. For example, the floating payments on a CMS swap might be made every 6 months at a rate equal to the 5-year swap rate. Usually there is a lag so that the payment on a particular payment date is equal to the swap rate observed on the previous payment date. Suppose that rates are set at times t_0, t_1, t_2, \dots , payments are made at times t_1, t_2, t_3, \dots , and L is the notional principal. The floating payment at time t_{i+1} is

$$\tau_i L S_i$$

where $\tau_i = t_{i+1} - t_i$ and S_i is the swap rate at time t_i .

Suppose that y_i is the forward value of the swap rate S_i . To value the payment at time t_{i+1} , it turns out to be correct to make a convexity/timing adjustment to the forward swap rate, so that the realized swap rate is assumed to be

$$y_i - \frac{1}{2} y_i^2 \sigma_{y,i}^2 \frac{G_i''(y_i)}{G_i'(y_i)} - \frac{y_i \tau_i F_i \rho_i \sigma_{y,i} \sigma_{F,i} t_i}{1 + F_i \tau_i} \quad (32.2)$$

rather than y_i . In this equation, $\sigma_{y,i}$ is the volatility of the forward swap rate, F_i is the

current forward interest rate between times t_i and t_{i+1} , $\sigma_{F,i}$ is the volatility of this forward rate, and ρ_i is the correlation between the forward swap rate and the forward interest rate. $G_i(x)$ is the price at time t_i of a bond as a function of its yield x . The bond pays coupons at rate y_i and has the same life and payment frequency as the swap from which the CMS rate is calculated. $G'_i(x)$ and $G''_i(x)$ are the first and second partial derivatives of G_i with respect to x . The volatilities $\sigma_{y,i}$ can be implied from swaptions; the volatilities $\sigma_{F,i}$ can be implied from caplet prices; the correlation ρ_i can be estimated from historical data.

Equation (32.2) involves a convexity and a timing adjustment. The term

$$-\frac{1}{2}y_i^2\sigma_{y,i}^2t_i\frac{G''_i(y_i)}{G'_i(y_i)}$$

is an adjustment similar the one in Example 29.2 of Section 29.1. It is based on the assumption that the swap rate S_i leads to only one payment at time t_i rather than to an annuity of payments. The term

$$-\frac{y_i\tau_iF_i\rho_i\sigma_{y,i}\sigma_{F,i}t_i}{1+F_i\tau_i}$$

is similar to the one in Section 29.2 and is an adjustment for the fact that the payment calculated from S_i is made at time t_{i+1} rather than t_i .

Example 32.3

In a 6-year CMS swap, the 5-year swap rate is received and a fixed rate of 5% is paid on a notional principal of \$100 million. The exchange of payments is semi-annual (both on the underlying 5-year swap and on the CMS swap itself). The exchange on a payment date is determined from the swap rate on the previous payment date. The term structure is flat at 5% per annum with semiannual compounding. All options on five-year swaps have a 15% implied volatility and all caplets with a 6-month tenor have a 20% implied volatility. The correlation between each cap rate and each swap rate is 0.7.

In this case, $y_i = 0.05$, $\sigma_{y,i} = 0.15$, $\tau_i = 0.5$, $F_i = 0.05$, $\sigma_{F,i} = 0.20$, and $\rho_i = 0.7$ for all i . Also,

$$G_i(x) = \sum_{i=1}^{10} \frac{2.5}{(1+x/2)^i} + \frac{100}{(1+x/2)^{10}}$$

so that $G'_i(y_i) = -437.603$ and $G''_i(y_i) = 2261.23$. Equation (32.2) gives the total convexity/timing adjustment as $0.0001197t_i$, or 1.197 basis points per year until the swap rate is observed. For example, for the purposes of valuing the CMS swap, the 5-year swap rate in 4 years' time should be assumed to be 5.0479% rather than 5% and the net cash flow received at the 4.5-year point should be assumed to be $0.5 \times 0.000479 \times 100,000,000 = \$23,940$. Other net cash flows are calculated similarly. Taking their present value, we find the value of the swap to be \$159,811.

A constant maturity Treasury swap (CMT swap) works similarly to a CMS swap except that the floating rate is the yield on a Treasury bond with a specified life. The analysis of a CMT swap is essentially the same as that for a CMS swap with S_i defined as the par yield on a Treasury bond with the specified life.

Differential Swaps

A *differential swap*, sometimes referred to as a *diff swap*, is an interest rate swap where a floating interest rate is observed in one currency and applied to a principal in another currency. Suppose that the LIBOR rate for the period between t_i and t_{i+1} in currency Y is applied to a principal in currency X with the payment taking place at time t_{i+1} . Define V_i as the forward interest rate between t_i and t_{i+1} in currency Y and W_i as the forward exchange rate for a contract with maturity t_{i+1} (expressed as the number of units of currency Y that equal one unit of currency X). If the LIBOR rate in currency Y were applied to a principal in currency Y, the cash flow at time t_{i+1} would be valued on the assumption that the LIBOR rate at time t_i equals V_i . From the analysis in Section 29.3, a quanto adjustment is necessary when it is applied to a principal in currency X. It is correct to value the cash flow on the assumption that the LIBOR rate equals

$$V_i + V_i \rho_i \sigma_{W,i} \sigma_{V,i} t_i \quad (32.3)$$

where $\sigma_{V,i}$ is the volatility of V_i , $\sigma_{W,i}$ is the volatility of W_i , and ρ_i is the correlation between V_i and W_i .

Example 32.4

Zero rates in both the US and Britain are flat at 5% per annum with annual compounding. In a 3-year diff swap agreement with annual payments, USD 12-month LIBOR is received and sterling 12-month LIBOR is paid with both being applied to a principal of 10 million pounds sterling. The volatility of all 1-year forward rates in the US is estimated to be 20%, the volatility of the forward USD/sterling exchange rate (dollars per pound) is 12% for all maturities, and the correlation between the two is 0.4.

In this case, $V_i = 0.05$, $\rho_i = 0.4$, $\sigma_{W,i} = 0.12$, $\sigma_{V,i} = 0.2$. The floating-rate cash flows dependent on the 1-year USD rate observed at time t_i should therefore be calculated on the assumption that the rate will be

$$0.05 + 0.05 \times 0.4 \times 0.12 \times 0.2 \times t_i = 0.05 + 0.00048t_i$$

This means that the net cash flows from the swap at times 1, 2, and 3 years should be assumed to be 0, 4,800, and 9,600 pounds sterling for the purposes of valuation. The value of the swap is therefore

$$\frac{0}{1.05} + \frac{4,800}{1.05^2} + \frac{9,600}{1.05^3} = 12,647$$

or 12,647 pounds sterling.

32.5 EQUITY SWAPS

In an equity swap, one party promises to pay the return on an equity index on a notional principal, while the other promises to pay a fixed or floating return on a notional principal. Equity swaps enable a fund managers to increase or reduce their exposure to an index without buying and selling stock. An equity swap is a convenient way of packaging a series of forward contracts on an index to meet the needs of the market.

Business Snapshot 32.3 Hypothetical Confirmation for an Equity Swap	
Trade date:	5-January, 2010
Effective date:	11-January, 2010
Business day convention (all dates):	Following business day
Holiday calendar:	US
Termination date:	11-January, 2015
<i>Equity amounts</i>	
Equity payer:	Microsoft
Equity principal:	USD 100 million
Equity index:	Total Return S&P 500 index
Equity payment:	$100(I_1 - I_0)/I_0$, where I_1 is the index level on the payment date and I_0 is the index level on the immediately preceding payment date. In the case of the first payment date, I_0 is the index level on 11-January, 2010
Equity payment dates:	Each 11-July and 11-January commencing 11-July, 2010, up to and including 11-January, 2015
<i>Floating amounts</i>	
Floating-rate payer:	Goldman Sachs
Floating-rate notional principal:	USD 100 million
Floating rate:	USD 6-month LIBOR
Floating-rate day count convention:	Actual/360
Floating-rate payment dates:	Each 11-July and 11-January commencing 11-July, 2010, up to and including 11-January, 2015

The equity index is usually a total return index where dividends are reinvested in the stocks comprising the index. An example of an equity swap is in Business Snapshot 32.3. In this, the 6-month return on the S&P 500 is exchanged for LIBOR. The principal on either side of the swap is \$100 million and payments are made every 6 months.

For an equity-for-floating swap such as that in Business Snapshot 32.3 the value at the start of its life is zero. This is because a financial institution can arrange to costlessly replicate the cash flows to one side by borrowing the principal on each payment date at LIBOR and investing it in the index until the next payment date with any dividends being reinvested. A similar argument shows that the swap is always worth zero immediately after a payment date.

Between payment dates the equity cash flow and the LIBOR cash flow at the next payment date must be valued. The LIBOR cash flow was fixed at the last reset date and so can be valued easily. The value of the equity cash flow is LE/E_0 , where L is the principal, E is the current value of the equity index, and E_0 is its value at the last payment date.³

³ See Technical Note 19 at www.rotman.utoronto.ca/~hull/TechnicalNotes for a more detailed discussion.

32.6 SWAPS WITH EMBEDDED OPTIONS

Some swaps contain embedded options. In this section we consider some commonly encountered examples.

Accrual Swaps

Accrual swaps are swaps where the interest on one side accrues only when the floating reference rate is within a certain range. Sometimes the range remains fixed during the entire life of the swap; sometimes it is reset periodically.

As a simple example of an accrual swap, consider a deal where a fixed rate Q is exchanged for 3-month LIBOR every quarter and the fixed rate accrues only on days when 3-month LIBOR is below 8% per annum. Suppose that the principal is L . In a normal swap the fixed-rate payer would pay QLn_1/n_2 on each payment date where n_1 is the number of days in the preceding quarter and n_2 is the number of days in the year. (This assumes that the day count is actual/actual.) In an accrual swap, this is changed to QLn_3/n_2 , where n_3 is the number of days in the preceding quarter that the 3-month LIBOR was below 8%. The fixed-rate payer saves QL/n_2 on each day when 3-month LIBOR is above 8%.⁴ The fixed-rate payer's position can therefore be considered equivalent to a regular swap plus a series of binary options, one for each day of the life of the swap. The binary options pay off QL/n_2 when the 3-month LIBOR is above 8%.

To generalize, suppose that the LIBOR cutoff rate (8% in the case just considered) is R_K and that payments are exchanged every τ years. Consider day i during the life of the swap and suppose that t_i is the time until day i . Suppose that the τ -year LIBOR rate on day i is R_i so that interest accrues when $R_i < R_K$. Define F_i as the forward value of R_i and σ_i as the volatility of F_i . (The latter is estimated from spot caplet volatilities.) Using the usual lognormal assumption, the probability that LIBOR is greater than R_K in a world that is forward risk neutral with respect to a zero-coupon bond maturing at time $t_i + \tau$ is $N(d_2)$, where

$$d_2 = \frac{\ln(F_i/R_K) - \sigma_i^2 t_i/2}{\sigma_i \sqrt{t_i}}$$

The payoff from the binary option is realized at the swap payment date following day i . Suppose that this is at time s_i . The probability that LIBOR is greater than R_K in a world that is forward risk neutral with respect to a zero-coupon bond maturing at time s_i is given by $N(d_2^*)$, where d_2^* is calculated using the same formula as d_2 , but with a small timing adjustment to F_i reflecting the difference between time $t_i + \tau$ and time s_i .

The value of the binary option corresponding to day i is

$$\frac{QL}{n_2} P(0, s_i) N(d_2^*)$$

The total value of the binary options is obtained by summing this expression for every

⁴ The usual convention is that, if a day is a holiday, the applicable rate is assumed to be the rate on the immediately preceding business day.

day in the life of the swap. The timing adjustment (causing d_2 to be replaced by d_2^*) is so small that, in practice, it is frequently ignored.

Cancelable Swap

A cancelable swap is a plain vanilla interest rate swap where one side has the option to terminate on one or more payment dates. Terminating a swap is the same as entering into the offsetting (opposite) swap. Consider a swap between Microsoft and Goldman Sachs. If Microsoft has the option to cancel, it can regard the swap as a regular swap plus a long position in an option to enter into the offsetting swap. If Goldman Sachs has the cancellation option, Microsoft has a regular swap plus a short position in an option to enter into the swap.

If there is only one termination date, a cancelable swap is the same as a regular swap plus a position in a European swaption. Consider, for example, a 10-year swap where Microsoft will receive 6% and pay LIBOR. Suppose that Microsoft has the option to terminate at the end of 6 years. The swap is a regular 10-year swap to receive 6% and pay LIBOR plus long position in a 6-year European option to enter into a 4-year swap where 6% is paid and LIBOR is received. (The latter is referred to as a 6×4 European swaption.) The standard market model for valuing European swaptions is described in Chapter 28.

When the swap can be terminated on a number of different payment dates, it is a regular swap plus a Bermudan-style swaption. Consider, for example, the situation where Microsoft has entered into a 5-year swap with semiannual payments where 6% is received and LIBOR is paid. Suppose that the counterparty has the option to terminate the swap on payment dates between year 2 and year 5. The swap is a regular swap plus a short position in a Bermudan-style swaption, where the Bermudan-style swaption is an option to enter into a swap that matures in 5 years and involves a fixed payment at 6% being received and a floating payment at LIBOR being paid. The swaption can be exercised on any payment date between year 2 and year 5. Methods for valuing Bermudan swaptions are discussed in Chapters 30 and 31.

Cancelable Compounding Swaps

Sometimes compounding swaps can be terminated on specified payment dates. On termination, the floating-rate payer pays the compounded value of the floating amounts up to the time of termination and the fixed-rate payer pays the compounded value of the fixed payments up to the time of termination.

Some tricks can be used to value cancelable compounding swaps. Suppose first that the floating rate is LIBOR and it is compounded at LIBOR. Assume that the principal amount of the swap is paid on both the fixed and floating sides of the swap at the end of its life. This is similar to moving from Table 7.1 to Table 7.2 for a vanilla swap. It does not change the value of the swap and has the effect of ensuring that the value of the floating side is always equals the notional principal on a payment date. To make the cancellation decision, we need only look at the fixed side. We construct an interest rate tree as outlined in Chapter 30. We roll back through the tree in the usual way valuing the fixed side. At each node where the swap can be canceled, we test whether it is optimal to keep the swap or cancel it. Canceling the swap in effect sets the fixed side equal to par. If we are paying fixed and receiving floating, our objective is to minimize

the value of the fixed side; if we are receiving fixed and paying floating, our objective is to maximize the value of the fixed side.

When the floating side is LIBOR plus a spread compounded at LIBOR, the cash flows corresponding to the spread rate of interest can be subtracted from the fixed side instead of adding them to the floating side. The option can then be valued as in the case where there is no spread.

When the compounding is at LIBOR plus a spread, an approximate approach is as follows:⁵

1. Calculate the value of the floating side of the swap at each cancellation date assuming forward rates are realized.
2. Calculate the value of the floating side of the swap at each cancellation date assuming that the floating rate is LIBOR and it is compounded at LIBOR.
3. Define the excess of step 1 over step 2 as the “value of spreads” on a cancellation date.
4. Treat the option in the way described above. In deciding whether to exercise the cancellation option, subtract the value of the spreads from the values calculated for the fixed side.

32.7 OTHER SWAPS

This chapter has discussed just a few of the swap structures in the market. In practice, the range of different contracts that trade is limited only by the imagination of financial engineers and the appetite of corporate treasurers for innovative risk management tools.

A swap that was very popular in the United States in the mid-1990s is an *index amortizing rate swap* (also called an *indexed principal swap*). In this, the principal reduces in a way dependent on the level of interest rates. The lower the interest rate, the greater the reduction in the principal. The fixed side of an indexed amortizing swap was originally designed to mirror approximately the return obtained by an investor on an agency mortgage-backed security after prepayment options are taken into account. The swap therefore exchanged the return on the mortgage-backed security for a floating-rate return.

Commodity swaps are now becoming increasingly popular. A company that consumes 100,000 barrels of oil per year could agree to pay \$8 million each year for the next 10 years and to receive in return 100,000 S , where S is the market price of oil per barrel. The agreement would in effect lock in the company’s oil cost at \$80 per barrel. An oil producer might agree to the opposite exchange, thereby locking in the price it realized for its oil at \$80 per barrel. Energy derivatives such as this will be discussed in Chapter 33.

A number of other types of swaps are discussed elsewhere in this book. For example, asset swaps are discussed in Chapter 23, total return swaps and various types of credit default swaps are covered in Chapter 24, and volatility and variance swaps are analyzed in Chapter 25.

Bizarre Deals

Some swaps have payoffs that are calculated in quite bizarre ways. An example is a deal entered into between Procter and Gamble and Bankers Trust in 1993 (see Business

⁵ This approach is not perfectly accurate in that it assumes that the decision to exercise the cancellation option is not influenced by future payments being compounded at a rate different from LIBOR.

Business Snapshot 32.4 Procter and Gamble's Bizarre Deal

A particularly bizarre swap is the so-called “5/30” swap entered into between Bankers Trust (BT) and Procter and Gamble (P&G) on November 2, 1993. This was a 5-year swap with semiannual payments. The notional principal was \$200 million. BT paid P&G 5.30% per annum. P&G paid BT the average 30-day CP (commercial paper) rate minus 75 basis points plus a spread. The average commercial paper rate was calculated by taking observations on the 30-day commercial paper rate each day during the preceding accrual period and averaging them.

The spread was zero for the first payment date (May 2, 1994). For the remaining nine payment dates, it was

$$\max \left[0, \frac{98.5 \left(\frac{5\text{-year CMT}\%}{5.78\%} \right) - (30\text{-year TSY price})}{100} \right]$$

In this, 5-year CMT is the constant maturity Treasury yield (i.e., the yield on a 5-year Treasury note, as reported by the US Federal Reserve). The 30-year TSY price is the midpoint of the bid and offer cash bond prices for the 6.25% Treasury bond maturing on August 2023. Note that the spread calculated from the formula is a decimal interest rate. It is not measured in basis points. If the formula gives 0.1 and the CP rate is 6%, the rate paid by P&G is 15.25%.

P&G were hoping that the spread would be zero and the deal would enable it to exchange fixed-rate funding at 5.30% for funding at 75 basis points less than the commercial paper rate. In fact, interest rates rose sharply in early 1994, bond prices fell, and the swap proved very, very expensive (see Problem 32.10).

Snapshot 32.4). The details of this transaction are in the public domain because it later became the subject of litigation.⁶

SUMMARY

Swaps have proved to be very versatile financial instruments. Many swaps can be valued by (a) assuming that LIBOR (or some other floating reference rate) will equal its forward value and (b) discounting the resulting cash flows at the LIBOR/swap rate. These include plain vanilla interest swaps, most types of currency swaps, swaps where the principal changes in a predetermined way, swaps where the payment dates are different on each side, and compounding swaps.

Some swaps require adjustments to the forward rates when they are valued. These adjustments are termed convexity, timing, or quanto adjustments. Among the swaps that require adjustments are LIBOR-in-arrears, CMS/CMT, and differential swaps.

Equity swaps involve the return on an equity index being exchanged for a fixed or floating rate of interest. They are usually designed so that they are worth zero immediately after a payment date, but they may have nonzero values between payment dates.

⁶ See D.J. Smith, “Aggressive Corporate Finance: A Close Look at the Procter and Gamble–Bankers Trust Leveraged Swap,” *Journal of Derivatives* 4, 4 (Summer 1997): 67–79.

Some swaps involve embedded options. An accrual swap is a regular swap plus a large portfolio of binary options (one for each day of the life of the swap). A cancelable swap is a regular swap plus a Bermudan swaption.

FURTHER READING

Chance, D., and Rich, D., “The Pricing of Equity Swap and Swaptions,” *Journal of Derivatives* 5, 4 (Summer 1998): 19–31.

Smith D.J., “Aggressive Corporate Finance: A Close Look at the Procter and Gamble–Bankers Trust Leveraged Swap,” *Journal of Derivatives*, 4, 4 (Summer 1997): 67–79.

Practice Questions (Answers in Solutions Manual)

- 32.1. Calculate all the fixed cash flows and their exact timing for the swap in Business Snapshot 32.1. Assume that the day count conventions are applied using target payment dates rather than actual payment dates.
- 32.2. Suppose that a swap specifies that a fixed rate is exchanged for twice the LIBOR rate. Can the swap be valued using the “assume forward rates are realized” rule?
- 32.3. What is the value of a 2-year fixed-for-floating compound swap where the principal is \$100 million and payments are made semiannually. Fixed interest is received and floating is paid? The fixed rate is 8% and it is compounded at 8.3% (both semiannually compounded). The floating rate is LIBOR plus 10 basis points and it is compounded at LIBOR plus 20 basis points. The LIBOR zero curve is flat at 8% with semiannual compounding.
- 32.4. What is the value of a 5-year swap where LIBOR is paid in the usual way and in return LIBOR compounded at LIBOR is received on the other side? The principal on both sides is \$100 million. Payment dates on the pay side and compounding dates on the receive side are every 6 months and the yield curve is flat at 5% with semiannual compounding.
- 32.5. Explain carefully why a bank might choose to discount cash flows on a currency swap at a rate slightly different from LIBOR.
- 32.6. Calculate the total convexity/timing adjustment in Example 32.3 of Section 32.4 if all cap volatilities are 18% instead of 20% and volatilities for all options on 5-year swaps are 13% instead of 15%. What should the 5-year swap rate in 3 years’ time be assumed for the purpose of valuing the swap? What is the value of the swap?
- 32.7. Explain why a plain vanilla interest rate swap and the compounding swap in Section 32.2 can be valued using the “assume forward rates are realized” rule, but a LIBOR-in-arrears swap in Section 32.4 cannot.
- 32.8. In the accrual swap discussed in the text, the fixed side accrues only when the floating reference rate lies below a certain level. Discuss how the analysis can be extended to cope with a situation where the fixed side accrues only when the floating reference rate is above one level and below another.

Further Questions

- 32.9. LIBOR zero rates are flat at 5% in the United States and flat at 10% in Australia (both annually compounded). In a 4-year swap Australian LIBOR is received and 9% is paid with both being applied to a USD principal of \$10 million. Payments are exchanged annually. The volatility of all 1-year forward rates in Australia is estimated to be 25%, the volatility of the forward USD/AUD exchange rate (AUD per USD) is 15% for all maturities, and the correlation between the two is 0.3. What is the value of the swap?
- 32.10. Estimate the interest rate paid by P&G on the 5/30 swap in Section 32.7 if (a) the CP rate is 6.5% and the Treasury yield curve is flat at 6% and (b) the CP rate is 7.5% and the Treasury yield curve is flat at 7% with semiannual compounding.
- 32.11. Suppose that you are trading a LIBOR-in-arrears swap with an unsophisticated counterparty who does not make convexity adjustments. To take advantage of the situation, should you be paying fixed or receiving fixed? How should you try to structure the swap as far as its life and payment frequencies?
- Consider the situation where the yield curve is flat at 10% per annum with annual compounding. All cap volatilities are 18%. Estimate the difference between the way a sophisticated trader and an unsophisticated trader would value a LIBOR-in-arrears swap where payments are made annually and the life of the swap is (a) 5 years, (b) 10 years, and (c) 20 years. Assume a notional principal of \$1 million.
- 32.12. Suppose that the LIBOR zero rate is flat at 5% with annual compounding. In a 5-year swap, company X pays a fixed rate of 6% and receives LIBOR. The volatility of the 2-year swap rate in 3 years is 20%.
- What is the value of the swap?
 - Use DerivaGem to calculate the value of the swap if company X has the option to cancel after 3 years.
 - Use DerivaGem to calculate the value of the swap if the counterparty has the option to cancel after 3 years.
 - What is the value of the swap if either side can cancel at the end of 3 years?

corn and wheat is the *stocks-to-use ratio*. This is the ratio of the year-end inventory to the year's usage. Typically it is between 20% and 40%. It has an impact on price volatility. As the ratio for a commodity becomes lower, the commodity's price becomes more sensitive to supply changes, so that the volatility increases.

There are reasons for supposing some level of mean reversion in agricultural prices. As prices decline, farmers find it less attractive to produce the commodity and supply decreases creating upward pressure on the price. Similarly, as the price of an agricultural commodity increases, farmers are more likely to devote resources to producing the commodity creating downward pressure on the price.

Prices of agricultural commodities tend to be seasonal, as storage is expensive and there is a limit to the length of time for which a product can be stored. Weather plays a key role in determining the price of many agricultural products. Frosts can decimate the Brazilian coffee crop, a hurricane in Florida is likely to have a big effect on the price of frozen orange juice, and so on. The volatility of the price of a commodity that is grown tends to be highest at pre-harvest times and then declines when the size of the crop is known. During the growing season, the price process for an agricultural commodity is liable to exhibit jumps because of the weather.

Many of the commodities that are grown and traded are used to feed livestock. (For example, the corn futures contract that is traded by the CME Group refers to the corn that is used to feed animals.) The price of livestock, and when slaughtering takes place, is liable to be dependent on the price of these commodities, which are in turn influenced by the weather.

33.2 METALS

Another important commodity category is metals. This includes gold, silver, platinum, palladium, copper, tin, lead, zinc, nickel, and aluminum. Metals have quite different characteristics from agricultural commodities. Their prices are unaffected by the weather and are not seasonal. They are extracted from the ground. They are divisible and are relatively easy to store. Some metals, such as copper, are used almost entirely in the manufacture of goods and should be classified as consumption assets. As explained in Section 5.1, others, such as gold and silver, are held purely for investment as well as for consumption and should be classified as investment assets.

As in the case of agricultural commodities, analysts monitor inventory levels to determine short-term price volatility. Exchange rate volatility may also contribute to volatility as the country where the metal is extracted is often different from the country in whose currency the price is quoted. In the long term, the price of a metal is determined by trends in the extent to which a metal is used in different production processes and new sources of the metal that are found. Changes in exploration and extraction methods, geopolitics, cartels, and environmental regulation also have an impact.

One potential source of supply for a metal is recycling. A metal might be used to create a product and, over the following 20 years, 10% of the metal might come back on the market as a result of a recycling process.

Metals that are investment assets are not usually assumed to follow mean-reverting processes because a mean-reverting process would give rise to an arbitrage opportunity for the investor. For metals that are consumption assets, there may be some mean

reversion. As the price of a metal increases, it is likely to become less attractive to use the metal in some production processes and more economically viable to extract the metal from difficult locations. As a result there will be downward pressure on the price. Similarly, as the price decreases, it is likely to become more attractive to use the metal in some production processes and less economically viable to extract the metal from difficult locations. As a result, there will be upward pressure on the price.

33.3 ENERGY PRODUCTS

Energy products are among the most important and actively traded commodities. A wide range of energy derivatives trade in both the over-the-counter market and on exchanges. Here we consider oil, natural gas, and electricity. There are reasons for supposing that all three follow mean reverting processes. As the price of a source of energy rises, it is likely to be consumed less and and produced more. This creates a downward pressure on prices. As the price of a source of energy declines, it is likely to be consumed more, but production is likely to be less economically viable. This creates upward pressure on the price.

Crude Oil

The crude oil market is the largest commodity market in the world, with global demand amounting to about 80 million barrels daily. Ten-year fixed-price supply contracts have been commonplace in the over-the-counter market for many years. These are swaps where oil at a fixed price is exchanged for oil at a floating price.

There are many grades of crude oil, reflecting variations in the gravity and the sulfur content. Two important benchmarks for pricing are Brent crude oil (which is sourced from the North Sea) and West Texas Intermediate (WTI) crude oil. Crude oil is refined into products such as gasoline, heating oil, fuel oil, and kerosene.

In the over-the-counter market, virtually any derivative that is available on common stocks or stock indices is now available with oil as the underlying asset. Swaps, forward contracts, and options are popular. Contracts sometimes require settlement in cash and sometimes require settlement by physical delivery (i.e., by delivery of oil).

Exchange-traded contracts are also popular. The CME Group and Intercontinental-Exchange (ICE) trade a number of oil futures and oil futures options contracts. Some of the futures contracts are settled in cash; others are settled by physical delivery. For example, the Brent crude oil futures traded on ICE have a cash settlement option; the light sweet crude oil futures traded on CME Group require physical delivery. In both cases, the amount of oil underlying one contract is 1,000 barrels. The CME Group also trades popular contracts on two refined products: heating oil and gasoline. In both cases, one contract is for the delivery of 42,000 gallons.

Natural Gas

The natural gas industry throughout the world went through a period of deregulation and the elimination of government monopolies in the 1980s and 1990s. The supplier of natural gas is now not necessarily the same company as the producer of the gas. Suppliers are faced with the problem of meeting daily demand.

A typical over-the-counter contract is for the delivery of a specified amount of natural gas at a roughly uniform rate over a 1-month period. Forward contracts, options, and swaps are available in the over-the-counter market. The seller of natural gas is usually responsible for moving the gas through pipelines to the specified location.

The CME Group trades a contract for the delivery of 10,000 million British thermal units of natural gas. The contract, if not closed out, requires physical delivery to be made during the delivery month at a roughly uniform rate to a particular hub in Louisiana. ICE trades a similar contract in London.

Natural gas is a popular source of energy for heating buildings. It is also used to produce electricity, which in turn is used for air-conditioning. As a result, demand for natural gas is seasonal and dependent on the weather.

Electricity

Electricity is an unusual commodity because it cannot easily be stored.¹ The maximum supply of electricity in a region at any moment is determined by the maximum capacity of all the electricity-producing plants in the region. In the United States there are 140 regions known as *control areas*. Demand and supply are first matched within a control area, and any excess power is sold to other control areas. It is this excess power that constitutes the wholesale market for electricity. The ability of one control area to sell power to another control area depends on the transmission capacity of the lines between the two areas. Transmission from one area to another involves a transmission cost, charged by the owner of the line, and there are generally some transmission or energy losses.

A major use of electricity is for air-conditioning systems. As a result the demand for electricity, and therefore its price, is much greater in the summer months than in the winter months. The nonstorability of electricity causes occasional very large movements in the spot price. Heat waves have been known to increase the spot price by as much as 1,000% for short periods of time.

Like natural gas, electricity has been through a period of deregulation and the elimination of government monopolies. This has been accompanied by the development of an electricity derivatives market. The CME Group now trades a futures contract on the price of electricity, and there is an active over-the-counter market in forward contracts, options, and swaps. A typical contract (exchange-traded or over-the-counter) allows one side to receive a specified number of megawatt hours for a specified price at a specified location during a particular month. In a 5×8 contract, power is received for five days a week (Monday to Friday) during the off-peak period (11 p.m. to 7 a.m.) for the specified month. In a 5×16 contract, power is received five days a week during the on-peak period (7 a.m. to 11 p.m.) for the specified month. In a 7×24 contract, it is received around the clock every day during the month. Option contracts have either daily exercise or monthly exercise. In the case of daily exercise, the option holder can choose on each day of the month (by giving one day's notice) whether to receive the specified amount of power at the specified strike price. When there is monthly exercise a

¹ Electricity producers with spare capacity sometimes use it to pump water to the top of their hydroelectric plants so that it can be used to produce electricity at a later time. This is the closest they can get to storing this commodity.

single decision on whether to receive power for the whole month at the specified strike price is made at the beginning of the month.

An interesting contract in electricity and natural gas markets is what is known as a *swing option* or *take-and-pay option*. In this contract a minimum and maximum for the amount of power that must be purchased at a certain price by the option holder is specified for each day during a month and for the month in total. The option holder can change (or swing) the rate at which the power is purchased during the month, but usually there is a limit on the total number of changes that can be made.

33.4 MODELING COMMODITY PRICES

To value derivatives, we are often interested in modeling the spot price of a commodity in the traditional risk-neutral world. From Section 17.7, the expected future price of the commodity in this world is the futures price.

A Simple Process

A simple process for a commodity price can be constructed by assuming that the expected growth rate in the commodity price is dependent solely on time and the volatility of the commodity price is constant. The risk-neutral process for the commodity price S then has the form

$$\frac{dS}{S} = \mu(t) dt + \sigma dz \quad (33.1)$$

and

$$F(t) = \hat{E}[S(t)] = S(0)e^{\int_0^t \mu(\tau) d\tau}$$

where $F(t)$ is the futures price for a contract with maturity t and \hat{E} denotes expected value in a risk-neutral world. It follows that

$$\ln F(t) = \ln S(0) + \int_0^t \mu(\tau) d\tau$$

Differentiating both sides with respect to time gives

$$\mu(t) = \frac{\partial}{\partial t} [\ln F(t)]$$

Example 33.1

Suppose that the futures prices of live cattle at the end of July 2008 are (in cents per pound) as follows:

August 2008	62.20
October 2008	60.60
December 2008	62.70
February 2009	63.37
April 2009	64.42
June 2009	64.40

These can be used to estimate the expected growth rate in live cattle prices in a risk-neutral world. For example, when the model in equation (33.1) is used, the

expected growth rate in live cattle prices between October and December 2008, in a risk-neutral world is

$$\ln\left(\frac{62.70}{60.60}\right) = 0.034$$

or 3.4% per 2 months with continuous compounding. On an annualized basis, this is 20.4% per annum.

Example 33.2

Suppose that the futures prices of live cattle are as in Example 33.1. A certain breeding decision would involve an investment of \$100,000 now and expenditures of \$20,000 in 3 months, 6 months, and 9 months. The result is expected to be that an extra cattle will be available for sale at the end of the year. There are two major uncertainties: the number of pounds of extra cattle that will be available for sale and the price per pound. The expected number of pounds is 300,000. The expected price of cattle in 1 year in a risk-neutral world is, from Example 33.1, 64.40 cents per pound. Assuming that the risk-free rate of interest is 10% per annum, the value of the investment (in thousands of dollars) is

$$-100 - 20e^{-0.1 \times 0.25} - 20e^{-0.1 \times 0.50} - 20e^{-0.1 \times 0.75} + 300 \times 0.644e^{-0.1 \times 1} = 17.729$$

This assumes that any uncertainty about the extra amount of cattle that will be available for sale has zero systematic risk and that there is no correlation between the amount of cattle that will be available for sale and the price.

Mean Reversion

As already discussed, most commodity prices follow mean-reverting processes. They tend to get pulled back to a central value. A more realistic process than equation (33.1) for the risk-neutral process followed by the commodity price S is

$$d \ln S = [\theta(t) - a \ln S] dt + \sigma dz \quad (33.2)$$

This incorporates mean reversion and is analogous to the lognormal process assumed for the short-term interest rate in Chapter 30. Note that this process is sometimes written

$$\frac{dS}{S} = [\theta^*(t) - a \ln S] dt + \sigma dz$$

From Itô's lemma, this is equivalent to the process in equation (33.2) when $\theta^*(t) = \theta(t) + \frac{1}{2}\sigma^2$.

The trinomial tree methodology in Section 30.7 can be adapted to construct a tree for S and determine the value of $\theta(t)$ in equation (33.2) such that $F(t) = \hat{E}[S(t)]$. We will illustrate the procedure by building a three-step tree for the situation where the current spot price is \$20 and the 1-year, 2-year, and 3-year futures prices are \$22, \$23, and \$24, respectively. Suppose that $a = 0.1$ and $\sigma = 0.2$ in equation (33.2). We first define a variable X that is initially zero and follows the process

$$dX = -aX dt + \sigma dz \quad (33.3)$$

Using the procedure in Section 30.7, a trinomial tree can be constructed for X . This is shown in Figure 33.1.

The variable $\ln S$ follows the same process as X except for a time-dependent drift. Analogously to Section 30.7, the tree for X can be converted to a tree for $\ln S$ by displacing the positions of nodes. This tree is shown in Figure 33.2. The initial node corresponds to a price of 20, so the displacement for that node is $\ln 20$. Suppose that the displacement of the nodes at 1 year is α_1 . The values of the X at the three nodes at the 1-year point are $+0.3464$, 0 , and -0.3464 . The corresponding values of $\ln S$ are $0.3464 + \alpha_1$, α_1 , and $-0.3464 + \alpha_1$. The values of S are therefore $e^{0.3464+\alpha_1}$, e^{α_1} , and $e^{-0.3464+\alpha_1}$, respectively. We require the expected value of S to equal the futures price. This means that

$$0.1667e^{0.3464+\alpha_1} + 0.6666e^{\alpha_1} + 0.1667e^{-0.3464+\alpha_1} = 22$$

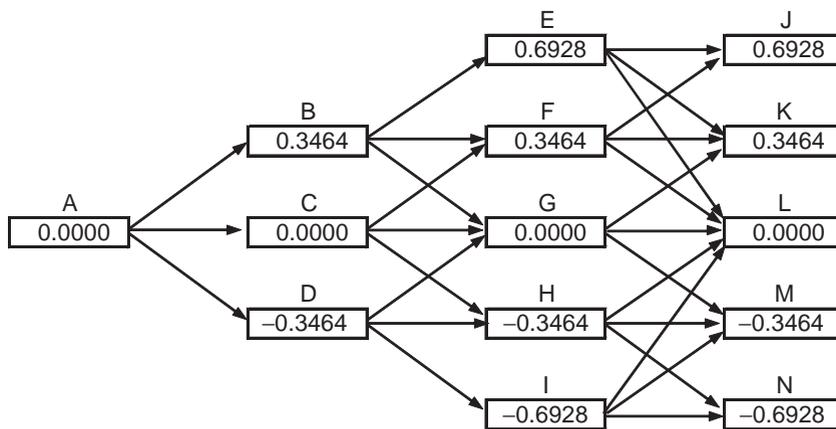
The solution to this is $\alpha_1 = 3.071$. The values of S at the 1-year point are therefore 30.49, 21.56, and 15.25.

At the 2-year point, we first calculate the probabilities of nodes E, F, G, H, and I being reached from the probabilities of nodes B, C, and D being reached. The probability of reaching node F is the probability of reaching node B times the probability of moving from B to F plus the probability of reaching node C times the probability of moving from C to F. This is

$$0.1667 \times 0.6566 + 0.6666 \times 0.1667 = 0.2206$$

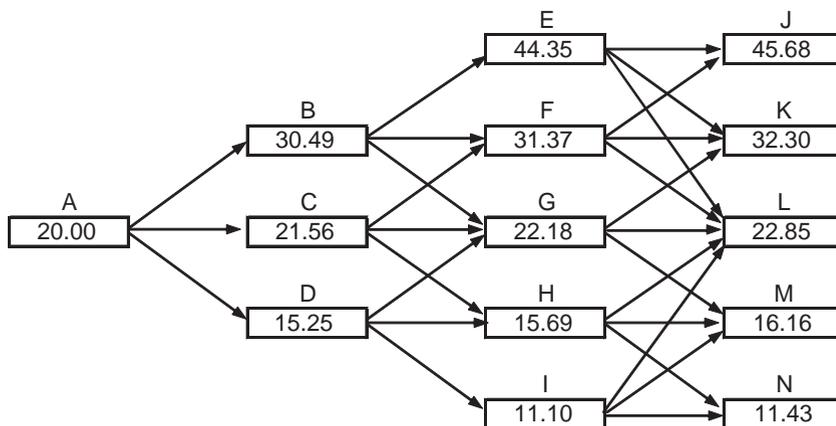
Similarly the probabilities of reaching nodes E, G, H, and I are 0.0203, 0.5183, 0.2206, and 0.0203, respectively. The amount α_2 by which the nodes at time 2 years are

Figure 33.1 Tree for X . Constructing this tree is the first stage in constructing a tree for the spot price of a commodity, S . Here p_u , p_m , and p_d are the probabilities of “up”, “middle”, and “down” movements from a node.



Node:	A	B	C	D	E	F	G	H	I
p_u :	0.1667	0.1217	0.1667	0.2217	0.8867	0.1217	0.1667	0.2217	0.0867
p_m :	0.6666	0.6566	0.6666	0.6566	0.0266	0.6566	0.6666	0.6566	0.0266
p_d :	0.1667	0.2217	0.1667	0.1217	0.0867	0.2217	0.1667	0.1217	0.8867

Figure 33.2 Tree for spot price of a commodity: p_u , p_m , and p_d are the probabilities of “up”, “middle”, and “down” movements from a node.



Node:	A	B	C	D	E	F	G	H	I
p_u :	0.1667	0.1217	0.1667	0.2217	0.8867	0.1217	0.1667	0.2217	0.0867
p_m :	0.6666	0.6566	0.6666	0.6566	0.0266	0.6566	0.6666	0.6566	0.0266
p_d :	0.1667	0.2217	0.1667	0.1217	0.0867	0.2217	0.1667	0.1217	0.8867

displaced must satisfy

$$0.0203e^{0.6928+\alpha_2} + 0.2206e^{0.3464+\alpha_2} + 0.5183e^{\alpha_2} + 0.2206e^{-0.3464+\alpha_2} + 0.0203e^{-0.6928+\alpha_2} = 23$$

The solution to this is $\alpha_2 = 3.099$. This means that the values of S at the 2-year point are 44.35, 31.37, 22.18, 15.69, and 11.10, respectively.

A similar calculation can be carried out at time 3 years. Figure 33.2 shows the resulting tree for S .

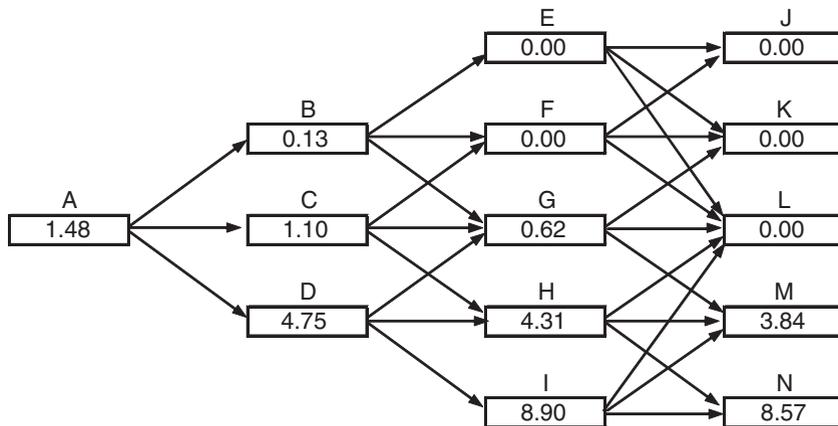
Example 33.3

Suppose that the tree in Figure 33.2 is used to price a 3-year American put option on the spot price of the commodity with a strike price of 20 when the interest rate (continuously compounded) is 3% per year. Rolling back through the tree in the usual way, we obtain Figure 33.3 showing that the value of the option is \$1.48. The option is exercised early at nodes D, H, and I. To obtain a more accurate value, a tree with many more time steps would be used. The futures prices would be interpolated to obtain futures prices for maturities corresponding to the end of every time step on this more detailed tree.

Interpolation and Seasonality

When a large number of time steps are used, it is necessary to interpolate between futures prices to obtain a futures price at the end of each time step. When there is seasonality, the interpolation procedure should reflect this. Suppose there are monthly

Figure 33.3 Valuation of an American put option with a strike price of \$20 using the tree in Figure 33.2.



Node:	A	B	C	D	E	F	G	H	I
p_u :	0.1667	0.1217	0.1667	0.2217	0.8867	0.1217	0.1667	0.2217	0.0867
p_m :	0.6666	0.6566	0.6666	0.6566	0.0266	0.6566	0.6666	0.6566	0.0266
p_d :	0.1667	0.2217	0.1667	0.1217	0.0867	0.2217	0.1667	0.1217	0.8867

time steps. One simple way of incorporating seasonality is to collect monthly historical data on the spot price and calculate the 12-month moving average of the price. A *percentage seasonal factor* can then be estimated as the average of the ratio of the spot price for the month to the 12-month moving average of spot prices that is centered (approximately) on the month.

The percentage seasonal factors are then used to deseasonalize the futures prices that are known. Monthly deseasonalized futures are then calculated using interpolation. These futures prices are then seasonalized using the percentage seasonal factors and the tree is built. Suppose, for example, that the futures prices are observed in the market for September and December as 40 and 44, respectively, and we want to calculate a futures prices for October and November. Suppose further that the percentage seasonality factors for September, October, November, and December are calculated from historical data as 0.95, 0.85, 0.8 and 1.1, respectively. The deseasonalized futures prices are $40/0.95 = 42.1$ for September and $44/1.1 = 40$ for December. The interpolated deseasonalized futures prices are 41.4 and 40.7 for October and November, respectively. The seasonalized futures prices that would be used in tree construction for October and November are $41.4 \times 0.85 = 35.2$ and $40.7 \times 0.8 = 32.6$, respectively.

As has been mentioned, the volatility of a commodity sometimes shows seasonality. For example, the prices of some agricultural commodities are more volatile during the growing season because of weather uncertainty. Volatility can be monitored using the methods discussed in Chapter 22, and a percentage seasonal factor for volatility can be estimated. The parameter σ can then be replaced by $\sigma(t)$ in equations (33.2) and (33.3). A procedure that can be used to construct a trinomial tree for the situation

where the volatility is a function of time is discussed in Technical Notes 9 and 16 at www.rotman.utoronto.ca/~hull/TechnicalNotes.

Jumps

Some commodities, such as electricity and natural gas, exhibit price jumps because of weather-related demand shocks. Other commodities, particularly those that are agricultural, are liable to exhibit price jumps because of weather-related supply shocks. Jumps can be incorporated into equation (33.2) so that the process for the spot price becomes

$$d \ln S = [\theta(t) - a \ln S] dt + \sigma dz + dp$$

where dp is the Poisson process generating the percentage jumps. This is similar to Merton's mixed jump-diffusion model for stock prices, which is described in Section 26.1. Once the jump frequency and jump size probability distribution have been chosen, the average increase in the commodity price at a future time t that is as a result of jumps can be calculated. To determine $\theta(t)$, the trinomial tree method can be used with the futures prices for maturity t reduced by this increase. Monte Carlo simulation can be used to implement the model, as explained in Sections 20.6 and 26.1.

Other Models

More-sophisticated models are sometimes used for oil prices. If y is the convenience yield, then the proportional drift of the spot price is $r - y$, where r is the short-term risk-free rate and a natural process to assume for the spot price is

$$\frac{dS}{S} = (r - y) dt + \sigma_1 dz_1$$

Gibson and Schwartz suggest that the convenience yield y be modeled as a mean-reverting process:²

$$dy = k(\alpha - y)dt + \sigma_2 dz_2$$

where k and α are constants and dz_2 is a Wiener process, which is correlated with the Wiener process dz_1 . To provide an exact fit to futures prices, α can be made a function of time.

Eydeland and Geman propose a stochastic volatility for gas and electricity prices.³ This is

$$\begin{aligned} \frac{dS}{S} &= a(b - \ln S) dt + \sqrt{V} dz_1 \\ dV &= c(d - V)dt + e\sqrt{V} dz_2 \end{aligned}$$

where a , b , c , d , and e are constants, and dz_1 and dz_2 are correlated Wiener processes. Later Geman proposed a model for oil where the reversion level b is also stochastic.⁴

² See R. Gibson and E.S. Schwartz, "Stochastic Convenience Yield and the Pricing of Oil Contingent Claims," *Journal of Finance*, 45 (1990): 959-76.

³ A. Eydeland and H. Geman, "Pricing Power Derivatives," *Risk*, September 1998.

⁴ H. Geman, "Scarcity and Price Volatility in Oil Markets," EDF Trading Technical Report, 2000.

33.5 WEATHER DERIVATIVES

Many companies are in the position where their performance is liable to be adversely affected by the weather.⁵ It makes sense for these companies to consider hedging their weather risk in much the same way as they hedge foreign exchange or interest rate risks.

The first over-the-counter weather derivatives were introduced in 1997. To understand how they work, we explain two variables:

HDD:	Heating degree days
CDD:	Cooling degree days

A day's HDD is defined as

$$\text{HDD} = \max(0, 65 - A)$$

and a day's CDD is defined as

$$\text{CDD} = \max(0, A - 65)$$

where A is the average of the highest and lowest temperature during the day at a specified weather station, measured in degrees Fahrenheit. For example, if the maximum temperature during a day (midnight to midnight) is 68° Fahrenheit and the minimum temperature is 44° Fahrenheit, $A = 56$. The daily HDD is then 9 and the daily CDD is 0.

A typical over-the-counter product is a forward or option contract providing a payoff dependent on the cumulative HDD or CDD during a month. For example, a derivatives dealer could in January 2011 sell a client a call option on the cumulative HDD during February 2012 at the Chicago O'Hare Airport weather station with a strike price of 700 and a payment rate of \$10,000 per degree day. If the actual cumulative HDD is 820, the payoff is \$1.2 million. Often contracts include a payment cap. If the payment cap in our example is \$1.5 million, the contract is the equivalent of a bull spread (see Chapter 11). The client has a long call option on cumulative HDD with a strike price of 700 and a short call option with a strike price of 850.

A day's HDD is a measure of the volume of energy required for heating during the day. A day's CDD is a measure of the volume of energy required for cooling during the day. Most weather derivative contracts are entered into by energy producers and consumers. But retailers, supermarket chains, food and drink manufacturers, health service companies, agriculture companies, and companies in the leisure industry are also potential users of weather derivatives. The Weather Risk Management Association (www.wrma.org) has been formed to serve the interests of the weather risk management industry.

In September 1999 the Chicago Mercantile Exchange (CME) began trading weather futures and European options on weather futures. The contracts are on the cumulative HDD and CDD for a month observed at a weather station. The contracts are settled in cash just after the end of the month once the HDD and CDD are known. One futures contract is on \$20 times the cumulative HDD or CDD for the month. The CME now offers weather futures and options on 42 cities throughout the world. It also offers futures and options on hurricanes, frost, and snowfall.

⁵ The US Department of Energy has estimated that one-seventh of the US economy is subject to weather risk.

33.6 INSURANCE DERIVATIVES

When derivative contracts are used for hedging purposes, they have many of the same characteristics as insurance contracts. Both types of contracts are designed to provide protection against adverse events. It is not surprising that many insurance companies have subsidiaries that trade derivatives and that many of the activities of insurance companies are becoming very similar to those of investment banks.

Traditionally the insurance industry has hedged its exposure to catastrophic (CAT) risks such as hurricanes and earthquakes using a practice known as reinsurance. Reinsurance contracts can take a number of forms. Suppose that an insurance company has an exposure of \$100 million to earthquakes in California and wants to limit this to \$30 million. One alternative is to enter into annual reinsurance contracts that cover on a pro rata basis 70% of its exposure. If California earthquake claims in a particular year total \$50 million, the costs to the company would then be only \$15 million. Another more popular alternative, involving lower reinsurance premiums, is to buy a series of reinsurance contracts covering what are known as *excess cost layers*. The first layer might provide indemnification for losses between \$30 million and \$40 million; the next layer might cover losses between \$40 million and \$50 million; and so on. Each reinsurance contract is known as an *excess-of-loss* reinsurance contract. The reinsurer has written a bull spread on the total losses. It is long a call option with a strike price equal to the lower end of the layer and short a call option with a strike price equal to the upper end of the layer.⁶

The principal providers of CAT reinsurance have traditionally been reinsurance companies and Lloyds syndicates (which are unlimited liability syndicates of wealthy individuals). In recent years the industry has come to the conclusion that its reinsurance needs have outstripped what can be provided from these traditional sources. It has searched for new ways in which capital markets can provide reinsurance. One of the events that caused the industry to rethink its practices was Hurricane Andrew in 1992, which caused about \$15 billion of insurance costs in Florida. This exceeded the total of relevant insurance premiums received in Florida during the previous seven years. If Hurricane Andrew had hit Miami, it is estimated that insured losses would have exceeded \$40 billion. Hurricane Andrew and other catastrophes have led to increases in insurance/reinsurance premiums.

The over-the-counter market has come up with a number of products that are alternatives to traditional reinsurance. The most popular is a CAT bond. This is a bond issued by a subsidiary of an insurance company that pays a higher-than-normal interest rate. In exchange for the extra interest the holder of the bond agrees to provide an excess-of-loss reinsurance contract. Depending on the terms of the CAT bond, the interest or principal (or both) can be used to meet claims. In the example considered above where an insurance company wants protection for California earthquake losses between \$30 million and \$40 million, the insurance company could issue CAT bonds with a total principal of \$10 million. In the event that the insurance company's California earthquake losses exceeded \$30 million, bondholders would lose some or all of their principal. As an alternative the insurance company could cover this excess cost layer by making a much bigger bond issue where only the bondholders' interest is at risk.

⁶ Reinsurance is also sometimes offered in the form of a lump sum if a certain loss level is reached. The reinsurer is then writing a cash-or-nothing binary call option on the losses.

33.7 PRICING WEATHER AND INSURANCE DERIVATIVES

One distinctive feature of weather and insurance derivatives is that there is no systematic risk (i.e., risk that is priced by the market) in their payoffs. This means that estimates made from historical data (real-world estimates) can also be assumed to apply to the risk-neutral world. Weather and insurance derivatives can therefore be priced by

1. Using historical data to estimate the expected payoff
2. Discounting the estimated expected payoff at the risk-free rate.

Another key feature of weather and insurance derivatives is the way uncertainty about the underlying variables grows with time. For a stock price, uncertainty grows roughly as the square root of time. Our uncertainty about a stock price in 4 years (as measured by the standard deviation of the logarithm of the price) is approximately twice that in 1 year. For a commodity price, mean reversion kicks in, but our uncertainty about a commodity's price in 4 years is still considerably greater than our uncertainty in 1 year. For weather, the growth of uncertainty with time is much less marked. Our uncertainty about the February HDD at a certain location in 4 years is usually only a little greater than our uncertainty about the February HDD at the same location in 1 year. Similarly, our uncertainty about earthquake losses for a period starting in 4 years is usually only a little greater than our uncertainty about earthquake losses for a similar period starting in 1 year.

Consider the valuation of an option on the cumulative HDD. We could collect 50 years of historical data and estimate a probability distribution for the HDD. This could be fitted to a lognormal or other probability distribution and the expected payoff on the option calculated. This would then be discounted at the risk-free rate to give the value of the option. The analysis could be refined by analyzing trends in the historical data and incorporating weather forecasts produced by meteorologists.

Example 33.4

Consider a call option on the cumulative HDD in February 2013 at the Chicago O'Hare Airport weather station with a strike price of 700 and a payment rate of \$10,000 per degree day. Suppose that the HDD is estimated from historical data to have a lognormal distribution with the mean HDD equal to 710 and the standard deviation of the natural logarithm of HDD equal to 0.07. From equation (14A.1), the expected payoff is

$$10,000 \times [710N(d_1) - 700N(d_2)]$$

where

$$d_1 = \frac{\ln(710/700) + 0.07^2/2}{0.07} = 0.2376$$

$$d_2 = \frac{\ln(710/700) - 0.07^2/2}{0.07} = 0.1676$$

or \$250,900. If the risk-free interest rate is 3% and the option is being valued in February 2012 (one year from maturity) the value of the option is

$$250,900 \times e^{-0.03 \times 1} = 243,400$$

or \$243,400.

We might want to adjust the the probability distribution of HDD for temperature trends. Suppose that a linear regression shows that the cumulative HDD for February is decreasing at the rate of 0.5 per year (perhaps because of global warming), so that the estimate of the mean HDD in February 2013 is only 697. Keeping the estimate of the standard deviation of the natural logarithm of the payoff the same, this would reduce the value of the expected payoff to \$180,400 and the value of the option to \$175,100.

Finally, suppose that long-range weather forecasters consider it likely that February 2013 will be particularly mild. The estimate of the expected HDD might then be reduced even further making the option even less valuable.

In the insurance area, Litzenberger *et al.* have shown that there are (as one would expect) no statistically significant correlation between the returns from CAT bonds and stock market returns.⁷ This confirms that there is no systematic risk and that valuations can be based on the actuarial data collected by insurance companies.

CAT bonds typically give a high probability of an above-normal rate of interest and a low probability of a big loss. Why would investors be interested in such instruments? The answer is that the expected return (taking account of possible losses) is higher than the return that can be earned on risk-free investments. However, the risk in CAT bonds can (at least in theory) be completely diversified away in a large portfolio. CAT bonds therefore have the potential to improve risk–return trade-offs.

33.8 HOW AN ENERGY PRODUCER CAN HEDGE RISKS

There are two components to the risks facing an energy producer. One is the risk associated with the market price for the energy (the price risk); the other is risk associated with the amount of energy that will be bought (the volume risk). Although prices do adjust to reflect volumes, there is a less-than-perfect relationship between the two, and energy producers have to take both into account when developing a hedging strategy. The price risk can be hedged using the energy derivative contracts. The volume risks can be hedged using the weather derivatives. Define:

Y : Profit for a month

P : Average energy prices for the month

T : Relevant temperature variable (HDD or CDD) for the month.

An energy producer can use historical data to obtain a best-fit linear regression relationship of the form

$$Y = a + bP + cT + \epsilon$$

where ϵ is the error term. The energy producer can then hedge risks for the month by taking a position of $-b$ in energy forwards or futures and a position of $-c$ in weather forwards or futures. The relationship can also be used to analyze the effectiveness of alternative option strategies.

⁷ R. H. Litzenberger, D. R. Beaglehole, and C. E. Reynolds, “Assessing Catastrophe Reinsurance-Linked Securities as a New Asset Class,” *Journal of Portfolio Management*, Winter 1996: 76–86.

SUMMARY

When there are risks to be managed, derivatives markets have been very innovative in developing products to meet the needs of the market.

There are a number of different types of commodity derivatives. The underlyings include agricultural products that are grown, livestock, metals, and energy products. The models used to value them usually incorporate mean reversion. Sometimes seasonality is modeled explicitly and jumps are incorporated. Energy derivatives with oil, natural gas, and electricity as the underlying are particularly important and have been the subject of models that are as sophisticated as the most sophisticated models used for stock prices, exchange rates, and interest rates.

In the weather derivatives market, two measures, HDD and CDD, have been developed to describe temperature during a month. These are used to define payoffs on both exchange-traded and over-the-counter derivatives. No doubt, as the weather derivatives market develops, contracts on rainfall, snow, and other weather-related variables will become more common.

Insurance derivatives are an alternative to traditional reinsurance as a way for insurance companies to manage the risk of a catastrophic event such as a hurricane or an earthquake. We may see other sorts of insurance, such as life and automobile insurance, being traded in a similar way in the future.

Weather and insurance derivatives have the property that the underlying variables have no systematic risk. This means that the derivatives can be valued by estimating expected payoffs using historical data and discounting the expected payoff at the risk-free rate.

FURTHER READING

On commodity derivatives

Clelow, L., and C. Strickland. *Energy Derivatives: Pricing and Risk Management*. Lacima Group, 2000.

Edwards, D.W. *Energy, Trading, and Investing: Trading, Risk Management and Structuring Deals in the Energy Markets*. Maidenhead: McGraw-Hill, 2010.

Eydeland, A., and H. Geman. "Pricing Power Derivatives," *Risk* (October 1998): 71–73.

Eydeland, A., and W. Krzysztof. *Energy and Power Risk Management*. Hoboken, NJ: Wiley, 2003.

Geman, H. *Commodities and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals, and Energy*. Chichester: Wiley, 2005.

Gibson, R., and E.S. Schwartz. "Stochastic Convenience Yield and the Pricing of Oil Contingent Claims," *Journal of Finance*, 45 (1990): 959–76.

On weather derivatives

Arditti, F., L. Cai, M. Cao, and R. McDonald. "Whether to Hedge," *RISK: Supplement on Weather Risk* (August 1999): 9–12.

Cao, M., and J. Wei. "Weather Derivatives Valuation and the Market Price of Weather Risk," *Journal of Futures Markets*, 24, 11 (November 2004), 1065–89.

Hunter, R. "Managing Mother Nature," *Derivatives Strategy* (February 1999).

On insurance derivatives

Canter, M. S., J. B. Cole, and R. L. Sandor. "Insurance Derivatives: A New Asset Class for the Capital Markets and a New Hedging Tool for the Insurance Industry," *Journal of Applied Corporate Finance* (Autumn 1997): 69–83.

Froot, K. A. "The Market for Catastrophe Risk: A Clinical Examination," *Journal of Financial Economics*, 60 (2001): 529–71.

Froot, K. A. *The Financing of Catastrophe Risk*. Chicago: University of Chicago Press, 1999.

Geman, H. "CAT Calls," *Risk* (September 1994): 86–89.

Hanley, M. "A Catastrophe Too Far," *Risk Supplement on Insurance* (July 1998).

Litzenberger, R. H., D. R. Beaglehole, and C. E. Reynolds. "Assessing Catastrophe Reinsurance-Linked Securities as a New Asset Class," *Journal of Portfolio Management* (Winter 1996): 76–86.

Practice Questions (Answers in Solutions Manual)

- 33.1. What is meant by HDD and CDD?
- 33.2. How is a typical natural gas forward contract structured?
- 33.3. Distinguish between the historical data and the risk-neutral approach to valuing a derivative. Under what circumstance do they give the same answer.
- 33.4. Suppose that each day during July the minimum temperature is 68° Fahrenheit and the maximum temperature is 82° Fahrenheit. What is the payoff from a call option on the cumulative CDD during July with a strike of 250 and a payment rate of \$5,000 per degree-day?
- 33.5. Why is the price of electricity more volatile than that of other energy sources?
- 33.6. Why is the historical data approach appropriate for pricing a weather derivatives contract and a CAT bond?
- 33.7. "HDD and CDD can be regarded as payoffs from options on temperature." Explain this statement.
- 33.8. Suppose that you have 50 years of temperature data at your disposal. Explain carefully the analyses you would carry out to value a forward contract on the cumulative CDD for a particular month.
- 33.9. Would you expect the volatility of the 1-year forward price of oil to be greater than or less than the volatility of the spot price? Explain your answer.
- 33.10. What are the characteristics of an energy source where the price has a very high volatility and a very high rate of mean reversion? Give an example of such an energy source.
- 33.11. How can an energy producer use derivatives markets to hedge risks?
- 33.12. Explain how a 5×8 option contract for May 2009 on electricity with daily exercise works. Explain how a 5×8 option contract for May 2009 on electricity with monthly exercise works. Which is worth more?
- 33.13. Explain how CAT bonds work.
- 33.14. Consider two bonds that have the same coupon, time to maturity, and price. One is a B-rated corporate bond. The other is a CAT bond. An analysis based on historical data shows that the expected losses on the two bonds in each year of their life is the same. Which bond would you advise a portfolio manager to buy and why?

Further Questions

- 33.15. An insurance company's losses of a particular type are to a reasonable approximation normally distributed with a mean of \$150 million and a standard deviation of \$50 million. (Assume no difference between losses in a risk-neutral world and losses in the real world.) The 1-year risk-free rate is 5%. Estimate the cost of the following:
- (a) A contract that will pay in 1 year's time 60% of the insurance company's costs on a pro rata basis
 - (b) A contract that pays \$100 million in 1 year's time if losses exceed \$200 million.
- 33.16. How is the tree in Figure 33.2 modified if the 1- and 2-year futures prices are \$21 and \$22 instead of \$22 and \$23, respectively. How does this affect the value of the American option in Example 33.3.

34

CHAPTER



Real Options

Up to now we have been almost entirely concerned with the valuation of financial assets. In this chapter we explore how the ideas we have developed can be extended to assess capital investment opportunities in real assets such as land, buildings, plant, and equipment. Often there are options embedded in these investment opportunities (the option to expand the investment, the option to abandon the investment, the option to defer the investment, and so on.) These options are very difficult to value using traditional capital investment appraisal techniques. The approach known as *real options* attempts to deal with this problem using option pricing theory.

The chapter starts by explaining the traditional approach to evaluating investments in real assets and shows how difficult it is to correctly value embedded options when this approach is used. It then explains how the risk-neutral valuation approach can be extended to handle the valuation of real assets and presents a number of examples illustrating the application of the approach in different situations.

34.1 CAPITAL INVESTMENT APPRAISAL

The traditional approach to valuing a potential capital investment project is known as the “net present value” (NPV) approach. The NPV of a project is the present value of its expected future incremental cash flows. The discount rate used to calculate the present value is a “risk-adjusted” discount rate, chosen to reflect the risk of the project. As the riskiness of the project increases, the discount rate also increases.

As an example, consider an investment that costs \$100 million and will last 5 years. The expected cash inflow in each year (in the real world) is estimated to be \$25 million. If the risk-adjusted discount rate is 12% (with continuous compounding), the net present value of the investment is (in millions of dollars)

$$-100 + 25e^{-0.12 \times 1} + 25e^{-0.12 \times 2} + 25e^{-0.12 \times 3} + 25e^{-0.12 \times 4} + 25e^{-0.12 \times 5} = -11.53$$

A negative NPV, such as the one we have just calculated, indicates that the project will reduce the value of the company to its shareholders and should not be undertaken. A positive NPV would indicate that the project should be undertaken because it will increase shareholder wealth.

The risk-adjusted discount rate should be the return required by the company, or the company's shareholders, on the investment. This can be calculated in a number of ways. One approach often recommended involves the capital asset pricing model (see the appendix to Chapter 3). The steps are as follows:

1. Take a sample of companies whose main line of business is the same as that of the project being contemplated.
2. Calculate the betas of the companies and average them to obtain a proxy beta for the project.
3. Set the required rate of return equal to the risk-free rate plus the proxy beta times the excess return of the market portfolio over the risk-free rate.

One problem with the traditional NPV approach is that many projects contain embedded options. Consider, for example, a company that is considering building a plant to manufacture a new product. Often the company has the option to abandon the project if things do not work out well. It may also have the option to expand the plant if demand for the output exceeds expectations. These options usually have quite different risk characteristics from the base project and require different discount rates.

To understand the problem here, return to the example at the beginning of Chapter 12. This involved a stock whose current price is \$20. In three months the price will be either \$22 or \$18. Risk-neutral valuation shows that the value of a three-month call option on the stock with a strike price of 21 is 0.633. Footnote 1 of Chapter 12 shows that if the expected return required by investors on the stock in the real world is 16% then the expected return required on the call option is 42.6%. A similar analysis shows that if the option is a put rather than a call the expected return required on the option is -52.5% . These analyses mean that if the traditional NPV approach were used to value the call option the correct discount rate would be 42.6%, and if it were used to value a put option the correct discount rate would be -52.5% . There is no easy way of estimating these discount rates. (We know them only because we are able to value the options another way.) Similarly, there is no easy way of estimating the risk-adjusted discount rates appropriate for cash flows when they arise from abandonment, expansion, and other options. This is the motivation for exploring whether the risk-neutral valuation principle can be applied to options on real assets as well as to options on financial assets.

Another problem with the traditional NPV approach lies in the estimation of the appropriate risk-adjusted discount rate for the base project (i.e., the project without embedded options). The companies that are used to estimate a proxy beta for the project in the three-step procedure above have expansion options and abandonment options of their own. Their betas reflect these options and may not therefore be appropriate for estimating a beta for the base project.

34.2 EXTENSION OF THE RISK-NEUTRAL VALUATION FRAMEWORK

In Section 27.1 the market price of risk for a variable θ was defined as

$$\lambda = \frac{\mu - r}{\sigma} \quad (34.1)$$

where r is the risk-free rate, μ is the return on a traded security dependent only on θ ,

and σ is its volatility. As shown in Section 27.1, the market price of risk, λ , does not depend on the particular traded security chosen.

Suppose that a real asset depends on several variables θ_i ($i = 1, 2, \dots$). Let m_i and s_i be the expected growth rate and volatility of θ_i so that

$$\frac{d\theta_i}{\theta_i} = m_i dt + s_i dz_i$$

where z_i is a Wiener process. Define λ_i as the market price of risk of θ_i . As explained in Section 27.9, risk-neutral valuation can be extended to show that any asset dependent on the θ_i can be valued by¹

1. Reducing the expected growth rate of each θ_i from m_i to $m_i - \lambda_i s_i$
2. Discounting cash flows at the risk-free rate.

Example 34.1

The cost of renting commercial real estate in a certain city is quoted as the amount that would be paid per square foot per year in a new 5-year rental agreement. The current cost is \$30 per square foot. The expected growth rate of the cost is 12% per annum, its volatility is 20% per annum, and its market price of risk is 0.3. A company has the opportunity to pay \$1 million now for the option to rent 100,000 square feet at \$35 per square foot for a 5-year period starting in 2 years. The risk-free rate is 5% per annum (assumed constant). Define V as the quoted cost per square foot of office space in 2 years. Assume that rent is paid annually in advance. The payoff from the option is

$$100,000A \max(V - 35, 0)$$

where A is an annuity factor given by

$$A = 1 + 1 \times e^{-0.05 \times 1} + 1 \times e^{-0.05 \times 2} + 1 \times e^{-0.05 \times 3} + 1 \times e^{-0.05 \times 4} = 4.5355$$

The expected payoff in a risk-neutral world is therefore

$$100,000 \times 4.5355 \times \hat{E}[\max(V - 35, 0)] = 453,550 \times \hat{E}[\max(V - 35, 0)]$$

where \hat{E} denotes expectations in a risk-neutral world. Using the result in equation (14A.1), this is

$$453,550[\hat{E}(V)N(d_1) - 35N(d_2)]$$

where

$$d_1 = \frac{\ln[\hat{E}(V)/35] + 0.2^2 \times 2/2}{0.2\sqrt{2}} \quad \text{and} \quad d_2 = \frac{\ln[\hat{E}(V)/35] - 0.2^2 \times 2/2}{0.2\sqrt{2}}$$

The expected growth rate in the cost of commercial real estate in a risk-neutral world is $m - \lambda s$, where m is the real-world growth rate, s is the volatility, and λ is the market price of risk. In this case, $m = 0.12$, $s = 0.2$, and $\lambda = 0.3$, so that the

¹ To see that this is consistent with regular risk-neutral valuation, suppose that θ_i is the price of a non-dividend-paying stock. Since this is the price of a traded security, equation (34.1) implies that $(m_i - r)/s_i = \lambda_i$, or $m_i - \lambda_i s_i = r$. The expected growth-rate adjustment is therefore the same as setting the return on the stock equal to the risk-free rate. For a proof of the more general result, see Technical Note 20 at:

expected risk-neutral growth rate is 0.06, or 6%, per year. It follows that $\hat{E}(V) = 30e^{0.06 \times 2} = 33.82$. Substituting this in the expression above gives the expected payoff in a risk-neutral world as \$1.5015 million. Discounting at the risk-free rate the value of the option is $1.5015e^{-0.05 \times 2} = \1.3586 million. This shows that it is worth paying \$1 million for the option.

34.3 ESTIMATING THE MARKET PRICE OF RISK

The real-options approach to evaluating an investment avoids the need to estimate risk-adjusted discount rates in the way described in Section 34.1, but it does require market price of risk parameters for all stochastic variables. When historical data are available for a particular variable, its market price of risk can be estimated using the capital asset pricing model. To show how this is done, we consider an investment asset dependent solely on the variable and define:

μ : Expected return of the investment asset

σ : Volatility of the return of the investment asset

λ : Market price of risk of the variable

ρ : Instantaneous correlation between the percentage changes in the variable and returns on a broad index of stock market prices

μ_m : Expected return on broad index of stock market prices

σ_m : Volatility of return on the broad index of stock market prices

r : Short-term risk-free rate

Because the investment asset is dependent solely on the market variable, the instantaneous correlation between its return and the broad index of stock market prices is also ρ . From the continuous-time version of the capital asset pricing model,

$$\mu - r = \frac{\rho\sigma}{\sigma_m}(\mu_m - r)$$

From equation (34.1), another expression for $\mu - r$ is

$$\mu - r = \lambda\sigma$$

It follows that

$$\lambda = \frac{\rho}{\sigma_m}(\mu_m - r) \quad (34.2)$$

This equation can be used to estimate λ .

Example 34.2

A historical analysis of company's sales, quarter by quarter, show that percentage changes in sales have a correlation of 0.3 with returns on the S&P 500 index. The volatility of the S&P 500 is 20% per annum and based on historical data the expected excess return of the S&P 500 over the risk-free rate is 5%. Equation (34.2) estimates the market price of risk for the company's sales as

$$\frac{0.3}{0.2} \times 0.05 = 0.075$$

When no historical data are available for the particular variable under consideration, other similar variables can sometimes be used as proxies. For example, if a plant is being constructed to manufacture a new product, data can be collected on the sales of other similar products. The correlation of the new product with the market index can then be assumed to be the same as that of these other products. In some cases, the estimate of ρ in equation (34.2) must be based on subjective judgment. If an analyst is convinced that a particular variable is unrelated to the performance of a market index, its market price of risk should be set to zero.

For some variables, it is not necessary to estimate the market price of risk because the process followed by a variable in a risk-neutral world can be estimated directly. For example, if the variable is the price of an investment asset, its total return in a risk-neutral world is the risk-free rate. If the variable is the short-term interest rate r , Chapter 30 shows how a risk-neutral process can be estimated from the initial term structure of interest rates.

For commodities, futures prices can be used to estimate the risk-neutral process, as discussed in Chapter 33. Example 33.2 provides a simple application of the real options approach by using futures prices to evaluate an investment involving the breeding of cattle.

34.4 APPLICATION TO THE VALUATION OF A BUSINESS

Traditional methods of business valuation, such as applying a price/earnings multiplier to current earnings, do not work well for new businesses. Typically a company's earnings are negative during its early years as it attempts to gain market share and establish relationships with customers. The company must be valued by estimating future earnings and cash flows under different scenarios.

The real options approach can be useful in this situation. A model relating the company's future cash flows to variables such as the sales growth rates, variable costs as a percent of sales, fixed costs, and so on, is developed. For key variables, a risk-neutral stochastic process is estimated as outlined in the previous two sections. A Monte Carlo simulation is then carried out to generate alternative scenarios for the net cash flows per year in a risk-neutral world. It is likely that under some of these scenarios the company does very well and under others it becomes bankrupt and ceases operations. (The simulation must have a built in rule for determining when bankruptcy happens.) The value of the company is the present value of the expected cash flow in each year using the risk-free rate for discounting. Business Snapshot 34.1 gives an example of the application of the approach to Amazon.com.

34.5 EVALUATING OPTIONS IN AN INVESTMENT OPPORTUNITY

As already mentioned, most investment projects involve options. These options can add considerable value to the project and are often either ignored or valued incorrectly. Examples of the options embedded in projects are:

1. *Abandonment options.* This is an option to sell or close down a project. It is an American put option on the project's value. The strike price of the option is the

Business Snapshot 34.1 Valuing Amazon.com

One of the earliest published attempts to value a company using the real options approach was Schwartz and Moon (2000), who considered Amazon.com at the end of 1999. They assumed the following stochastic processes for the company's sales revenue R and its revenue growth rate μ :

$$\frac{dR}{R} = \mu dt + \sigma(t) dz_1$$

$$d\mu = \kappa(\bar{\mu} - \mu) dt + \eta(t) dz_2$$

They assumed that the two Wiener processes dz_1 and dz_2 were uncorrelated and made reasonable assumptions about $\sigma(t)$, $\eta(t)$, κ , and $\bar{\mu}$ based on available data.

They assumed the cost of goods sold would be 75% of sales, other variable expenses would be 19% of sales, and fixed expenses would be \$75 million per quarter. The initial sales level was \$356 million, the initial tax loss carry forward was \$559 million, and the tax rate was assumed to be 35%. The market price of risk for R was estimated from historical data using the approach described in the previous section. The market price of risk for μ was assumed to be zero.

The time horizon for the analysis was 25 years and the terminal value of the company was assumed to be ten times pretax operating profit. The initial cash position was \$906 million and the company was assumed to go bankrupt if the cash balance became negative.

Different future scenarios were generated in a risk-neutral world using Monte Carlo simulation. The evaluation of the scenarios involved taking account of the possible exercise of convertible bonds and the possible exercise of employee stock options. The value of the company to the share holders was calculated as the present value of the net cash flows discounted at the risk-free rate.

Using these assumptions, Schwartz and Moon provided an estimate of the value of Amazon.com's shares at the end of 1999 equal to \$12.42. The market price at the time was \$76.125 (although it declined sharply in 2000). One of the key advantages of the real-options approach is that it identifies the key assumptions. Schwartz and Moon found that the estimated share value was very sensitive to $\eta(t)$, the volatility of the growth rate. This was an important source of optionality. A small increase in $\eta(t)$ leads to more optionality and a big increase in the value of Amazon.com shares.

liquidation (or resale) value of the project less any closing-down costs. When the liquidation value is low, the strike price can be negative. Abandonment options mitigate the impact of very poor investment outcomes and increase the initial valuation of a project.

2. *Expansion options.* This is the option to make further investments and increase the output if conditions are favorable. It is an American call option on the value of additional capacity. The strike price of the call option is the cost of creating this additional capacity discounted to the time of option exercise. The strike price often depends on the initial investment. If management initially choose to build capacity in excess of the expected level of output, the strike price can be relatively small.

3. *Contraction options.* This is the option to reduce the scale of a project's operation. It is an American put option on the value of the lost capacity. The strike price is the present value of the future expenditures saved as seen at the time of exercise of the option.
4. *Options to defer.* One of the most important options open to a manager is the option to defer a project. This is an American call option on the value of the project.
5. *Options to extend.* Sometimes it is possible to extend the life of an asset by paying a fixed amount. This is a European call option on the asset's future value.

Example

As an example of the evaluation of an investment with embedded options, consider a company that has to decide whether to invest \$15 million to extract 6 million units of a commodity from a certain source at the rate of 2 million units per year for 3 years. The fixed costs of operating the equipment are \$6 million per year and the variable costs are \$17 per unit of the commodity extracted. We assume that the risk-free interest rate is 10% per annum for all maturities, that the spot price of the commodity is \$20, and that the 1-, 2-, and 3-year futures prices are \$22, \$23, and \$24, respectively.

Evaluation with No Embedded Options

First consider the case where the project has no embedded options. The expected prices of the commodity in 1, 2, and 3 years' time in a risk-neutral world are \$22, \$23, and \$24, respectively. The expected payoff from the project (in millions of dollars) in a risk-neutral world can be calculated from the cost data as 4.0, 6.0, and 8.0 in years 1, 2, and 3, respectively. The value of the project is therefore

$$-15.0 + 4.0e^{-0.1 \times 1} + 6.0e^{-0.1 \times 2} + 8.0e^{0.1 \times 3} = -0.54$$

This analysis indicates that the project should not be undertaken because it would reduce shareholder wealth by 0.54 million.

Use of a Tree

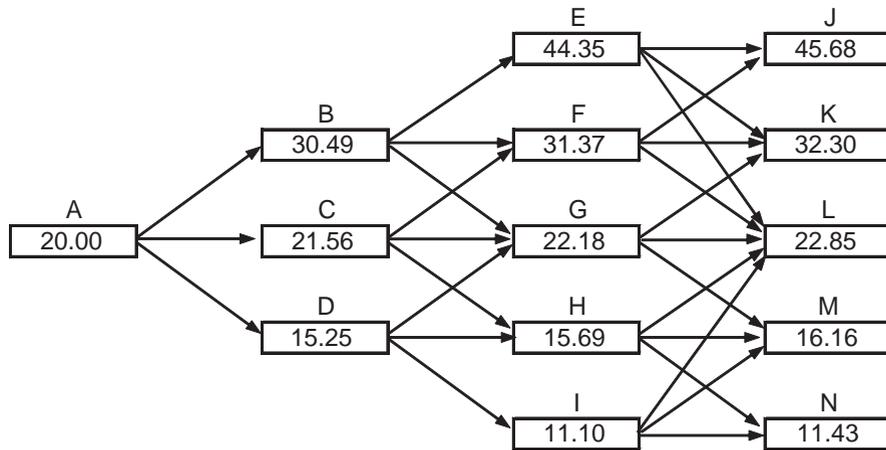
We now assume that the spot price of the commodity follows the process

$$d \ln S = [\theta(t) - a \ln S] dt + \sigma dz \quad (34.3)$$

where $a = 0.1$ and $\sigma = 0.2$. In Section 33.4, we showed how a tree can be constructed for commodity prices using the same example as the one considered here. The tree is shown in Figure 34.1 (which is the same as Figure 33.2). The process represented by the tree is consistent with the process assumed for S , the assumed values of a and σ , and the assumed 1-, 2-, and 3-year futures prices.

We do not need to use a tree to value the project when there are no embedded options. (We have already shown that the base value of the project without options is -0.54 .) However, before we move on to consider options, it will be instructive, as well as useful for future calculations, for us to use the tree to value the project in the absence of embedded options and verify that we get the same answer as that obtained earlier.

Figure 34.1 Tree for spot price of a commodity: p_u , p_m , and p_d are the probabilities of “up”, “middle”, and “down” movements from a node.



Node:	A	B	C	D	E	F	G	H	I
p_u :	0.1667	0.1217	0.1667	0.2217	0.8867	0.1217	0.1667	0.2217	0.0867
p_m :	0.6666	0.6566	0.6666	0.6566	0.0266	0.6566	0.6666	0.6566	0.0266
p_d :	0.1667	0.2217	0.1667	0.1217	0.0867	0.2217	0.1667	0.1217	0.8867

Figure 34.2 shows the value of the project at each node of Figure 34.1. Consider, for example, node H. There is a 0.2217 probability that the commodity price at the end of the third year is 22.85, so that the third-year profit is $2 \times 22.85 - 2 \times 17 - 6 = 5.70$. Similarly, there is a 0.6566 probability that the commodity price at the end of the third year is 16.16, so that the profit is -7.68 and there is a 0.1217 probability that the commodity price at the end of the third year is 11.43, so that the profit is -17.14 . The value of the project at node H in Figure 34.2 is therefore

$$[0.2217 \times 5.70 + 0.6566 \times (-7.68) + 0.1217 \times (-17.14)]e^{-0.1 \times 1} = -5.31$$

As another example, consider node C. There is a 0.1667 chance of moving to node F where the commodity price is 31.37. The second year cash flow is then

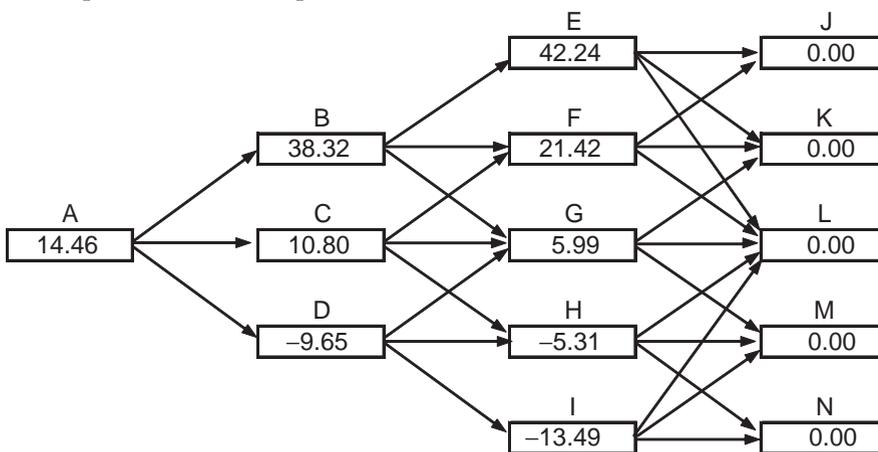
$$2 \times 31.37 - 2 \times 17 - 6 = 22.74$$

The value of subsequent cash flows at node F is 21.42. The total value of the project if we move to node F is therefore $21.42 + 22.74 = 44.16$. Similarly the total value of the project if we move to nodes G and H are 10.35 and -13.93 , respectively. The value of the project at node C is therefore

$$[0.1667 \times 44.16 + 0.6666 \times 10.35 + 0.1667 \times (-13.93)]e^{-0.1 \times 1} = 10.80$$

Figure 34.2 shows that the value of the project at the initial node A is 14.46. When the initial investment is taken into account the value of the project is therefore -0.54 . This is in agreement with our earlier calculations.

Figure 34.2 Valuation of base project with no embedded options: p_u , p_m , and p_d are the probabilities of “up”, “middle”, and “down” movements from a node.



Node:	A	B	C	D	E	F	G	H	I
p_u :	0.1667	0.1217	0.1667	0.2217	0.8867	0.1217	0.1667	0.2217	0.0867
p_m :	0.6666	0.6566	0.6666	0.6566	0.0266	0.6566	0.6666	0.6566	0.0266
p_d :	0.1667	0.2217	0.1667	0.1217	0.0867	0.2217	0.1667	0.1217	0.8867

Option to Abandon

Suppose now that the company has the option to abandon the project at any time. We suppose that there is no salvage value and no further payments are required once the project has been abandoned. Abandonment is an American put option with a strike price of zero and is valued in Figure 34.3. The put option should not be exercised at nodes E, F, and G because the value of the project is positive at these nodes. It should be exercised at nodes H and I. The value of the put option is 5.31 and 13.49 at nodes H and I, respectively. Rolling back through the tree, the value of the abandonment put option at node D if it is not exercised is

$$(0.1217 \times 13.49 + 0.6566 \times 5.31 + 0.2217 \times 0)e^{-0.1 \times 1} = 4.64$$

The value of exercising the put option at node D is 9.65. This is greater than 4.64, and so the put should be exercised at node D. The value of the put option at node C is

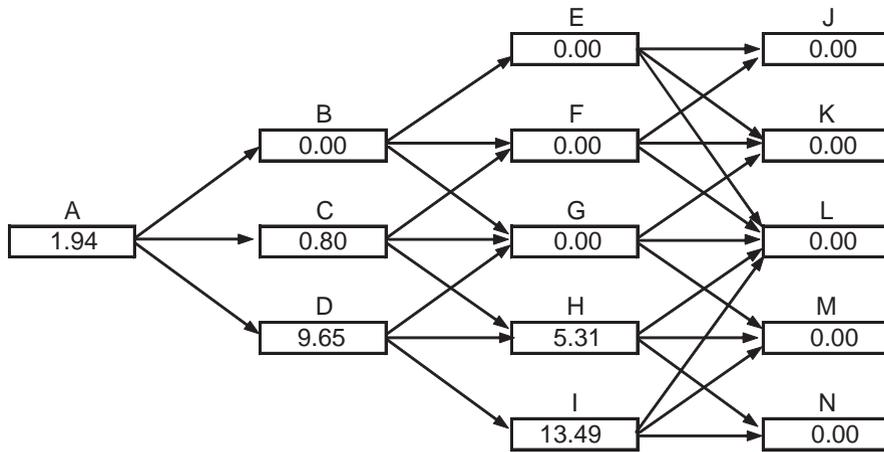
$$[0.1667 \times 0 + 0.6666 \times 0 + 0.1667 \times (5.31)]e^{-0.1 \times 1} = 0.80$$

and the value at node A is

$$(0.1667 \times 0 + 0.6666 \times 0.80 + 0.1667 \times 9.65)e^{-0.1 \times 1} = 1.94$$

The abandonment option is therefore worth \$1.94 million. It increases the value of the project from $-\$0.54$ million to $+\$1.40$ million. A project that was previously unattractive now has a positive value to shareholders.

Figure 34.3 Valuation of option to abandon the project: p_u , p_m , and p_d are the probabilities of “up”, “middle”, and “down” movements from a node.



Node:	A	B	C	D	E	F	G	H	I
p_u :	0.1667	0.1217	0.1667	0.2217	0.8867	0.1217	0.1667	0.2217	0.0867
p_m :	0.6666	0.6566	0.6666	0.6566	0.0266	0.6566	0.6666	0.6566	0.0266
p_d :	0.1667	0.2217	0.1667	0.1217	0.0867	0.2217	0.1667	0.1217	0.8867

Option to Expand

Suppose next that the company has no abandonment option. Instead it has the option at any time to increase the scale of the project by 20%. The cost of doing this is \$2 million. Production increases from 2.0 to 2.4 million units per year. Variable costs remain \$17 per unit and fixed costs increase by 20% from \$6.0 million to \$7.2 million. This is an American call option to buy 20% of the base project in Figure 34.2 for \$2 million. The option is valued in Figure 34.4. At node E, the option should be exercised. The payoff is $0.2 \times 42.24 - 2 = 6.45$. At node F, it should also be exercised for a payoff of $0.2 \times 21.42 - 2 = 2.28$. At nodes G, H, and I, the option should not be exercised. At node B, exercising is worth more than waiting and the option is worth $0.2 \times 38.32 - 2 = 5.66$. At node C, if the option is not exercised, it is worth

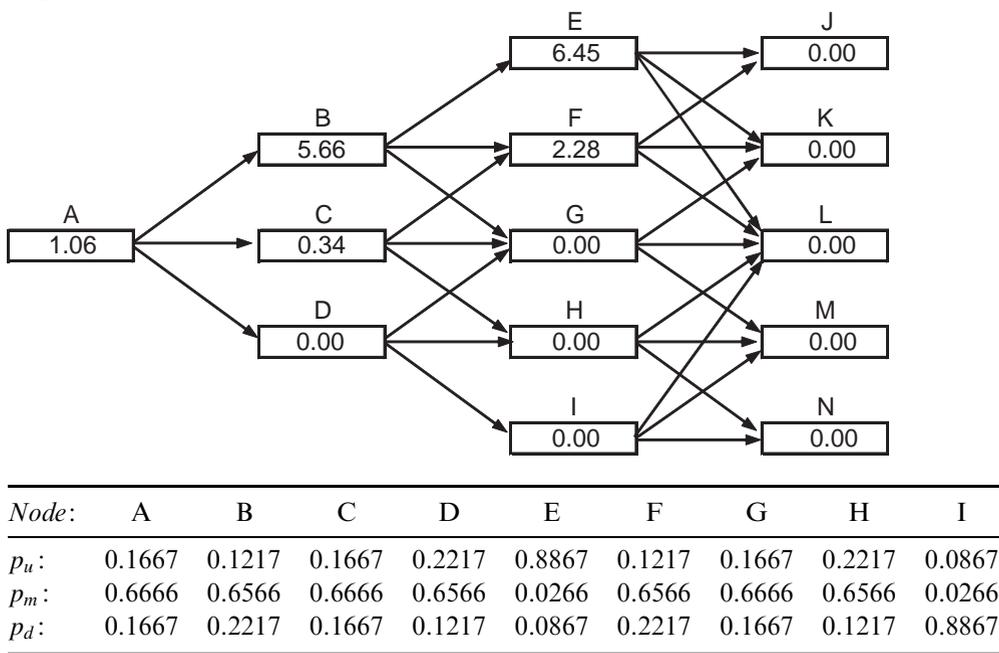
$$(0.1667 \times 2.28 + 0.6666 \times 0.00 + 0.1667 \times 0.00)e^{-0.1 \times 1} = 0.34$$

If the option is exercised, it is worth $0.2 \times 10.80 - 2 = 0.16$. The option should therefore not be exercised at node C. At node A, if not exercised, the option is worth

$$(0.1667 \times 5.66 + 0.6666 \times 0.34 + 0.1667 \times 0.00)e^{-0.1 \times 1} = 1.06$$

If the option is exercised it is worth $0.2 \times 14.46 - 2 = 0.89$. Early exercise is therefore not optimal at node A. In this case, the option increases the value of the project from -0.54 to $+0.52$. Again we find that a project that previously had a negative value now has a positive value.

Figure 34.4 Valuation of option to expand the project: p_u , p_m , and p_d are the probabilities of “up”, “middle”, and “down” movements from a node.



The expansion option in Figure 34.4 is relatively easy to value because, once the option has been exercised, all subsequent cash inflows and outflows increase by 20%. In the case where fixed costs remain the same or increase by less than 20%, it is necessary to keep track of more information at the nodes of Figure 34.4. Specifically, we need to record the following in order to calculate the payoff from exercising the option:

1. The present value of subsequent fixed costs
2. The present value of subsequent revenues net of variable costs.

Multiple Options

When a project has two or more options, they are typically not independent. The value of having both option A and option B, for example, is generally not the sum of the values of the two options. To illustrate this, suppose that the company we have been considering has both abandonment and expansion options. The project cannot be expanded if it has already been abandoned. Moreover, the value of the put option to abandon depends on whether the project has been expanded.²

These interactions between the options in our example can be handled by defining four states at each node:

1. Not already abandoned; not already expanded
2. Not already abandoned; already expanded

² As it happens, the two options do not interact in Figures 34.3 and 34.4. However, the interactions between the options would become an issue if a larger tree with smaller time steps were built.

3. Already abandoned; not already expanded
4. Already abandoned; already expanded.

When we roll back through the tree we calculate the combined value of the options at each node for all four alternatives. This approach to valuing path-dependent options is discussed in more detail in Section 26.5.

Several Stochastic Variables

When there are several stochastic variables, the value of the base project is usually determined by Monte Carlo simulation. The valuation of the project's embedded options is then more difficult because a Monte Carlo simulation works from the beginning to the end of a project. When we reach a certain point, we do not have information on the present value of the project's future cash flows. However, the techniques mentioned in Section 26.8 for valuing American options using Monte Carlo simulation can sometimes be used.

As an illustration of this point, Schwartz and Moon (2000) explain how their Amazon.com analysis outlined in Business Snapshot 34.1 could be extended to take account of the option to abandon (i.e. the option to declare bankruptcy) when the value of future cash flows is negative.³ At each time step, a polynomial relationship between the value of not abandoning and variables such as the current revenue, revenue growth rate, volatilities, cash balances, and loss carry forwards is assumed. Each simulation trial provides an observation for obtaining a least-squares estimate of the relationship at each time. This is the Longstaff and Schwartz approach of Section 26.8.⁴

SUMMARY

This chapter has investigated how the ideas developed earlier in the book can be applied to the valuation of real assets and options on real assets. It has shown how the risk-neutral valuation principle can be used to value a project dependent on any set of variables. The expected growth rate of each variable is adjusted to reflect its market price of risk. The value of the asset is then the present value of its expected cash flows discounted at the risk-free rate.

Risk-neutral valuation provides an internally consistent approach to capital investment appraisal. It also makes it possible to value the options that are embedded in many of the projects that are encountered in practice. This chapter has illustrated the approach by applying it to the valuation of Amazon.com at the end of 1999 and the valuation of a project involving the extraction of a commodity.

FURTHER READING

Amran, M., and N. Kulatilaka, *Real Options*, Boston, MA: Harvard Business School Press, 1999.

³ The analysis in Section 34.4 assumed that bankruptcy occurs when the cash balance falls below zero, but this is not necessarily optimal for Amazon.com.

⁴ F. A. Longstaff and E. S. Schwartz, "Valuing American Options by Simulation: A Simple Least-Squares Approach," *Review of Financial Studies*, 14, 1 (Spring 2001): 113–47.

Copeland, T., T. Koller, and J. Murrin, *Valuation: Measuring and Managing the Value of Companies*, 3rd edn. New York: Wiley, 2000.

Copeland, T., and V. Antikarov, *Real Options: A Practitioners Guide*, New York: Texere, 2003.

Schwartz, E.S., and M. Moon, "Rational Pricing of Internet Companies," *Financial Analysts Journal*, May/June (2000): 62–75.

Trigeorgis, L., *Real Options: Managerial Flexibility and Strategy in Resource Allocation*, Cambridge, MA: MIT Press, 1996.

Practice Questions (Answers in Solutions Manual)

- 34.1. Explain the difference between the net present value approach and the risk-neutral valuation approach for valuing a new capital investment opportunity. What are the advantages of the risk-neutral valuation approach for valuing real options?
- 34.2. The market price of risk for copper is 0.5, the volatility of copper prices is 20% per annum, the spot price is 80 cents per pound, and the 6-month futures price is 75 cents per pound. What is the expected percentage growth rate in copper prices over the next 6 months?
- 34.3. Consider a commodity with constant volatility σ and an expected growth rate that is a function solely of time. Show that, in the traditional risk-neutral world,

$$\ln S_T \sim \phi[(\ln F(T) - \frac{1}{2}\sigma^2 T, \sigma^2 T]$$

where S_T is the value of the commodity at time T , $F(t)$ is the futures price at time 0 for a contract maturing at time t , and $\phi(m, v)$ is a normal distribution with mean m and variance v .

- 34.4. Derive a relationship between the convenience yield of a commodity and its market price of risk.
- 34.5. The correlation between a company's gross revenue and the market index is 0.2. The excess return of the market over the risk-free rate is 6% and the volatility of the market index is 18%. What is the market price of risk for the company's revenue?
- 34.6. A company can buy an option for the delivery of 1 million units of a commodity in 3 years at \$25 per unit. The 3-year futures price is \$24. The risk-free interest rate is 5% per annum with continuous compounding and the volatility of the futures price is 20% per annum. How much is the option worth?
- 34.7. A driver entering into a car lease agreement can obtain the right to buy the car in 4 years for \$10,000. The current value of the car is \$30,000. The value of the car, S , is expected to follow the process $dS = \mu S dt + \sigma S dz$, where $\mu = -0.25$, $\sigma = 0.15$, and dz is a Wiener process. The market price of risk for the car price is estimated to be -0.1 . What is the value of the option? Assume that the risk-free rate for all maturities is 6%.

Further Questions

- 34.8. Suppose that the spot price, 6-month futures price, and 12-month futures price for wheat are 250, 260, and 270 cents per bushel, respectively. Suppose that the price of wheat follows the process in equation (34.3) with $a = 0.05$ and $\sigma = 0.15$. Construct a two-time-step tree for the price of wheat in a risk-neutral world.

A farmer has a project that involves an expenditure of \$10,000 and a further expenditure of \$90,000 in 6 months. It will increase wheat that is harvested and sold by 40,000 bushels in 1 year. What is the value of the project? Suppose that the farmer can abandon the project in 6 months and avoid paying the \$90,000 cost at that time. What is the value of the abandonment option? Assume a risk-free rate of 5% with continuous compounding.

34.9. In the example considered in Section 34.5:

- (a) What is the value of the abandonment option if it costs \$3 million rather than zero?
- (b) What is the value of the expansion option if it costs \$5 million rather than \$2 million?



35

CHAPTER

Derivatives Mishaps and What We Can Learn from Them

Since the mid-1980s there have been some spectacular losses in derivatives markets. The biggest losses have come from the trading of products created from residential mortgages in the US and were discussed in Chapter 8. Some of the other losses made by financial institutions are listed in Business Snapshot 35.1, and some of those made by nonfinancial organizations in Business Snapshot 35.2. What is remarkable about these lists is the number of situations where huge losses arose from the activities of a single employee. In 1995, Nick Leeson's trading brought a 200-year-old British bank, Barings, to its knees; in 1994, Robert Citron's trading led to Orange County, a municipality in California, losing about \$2 billion. Joseph Jett's trading for Kidder Peabody lost \$350 million. John Rusnak's losses of \$700 million for Allied Irish Bank came to light in 2002. In 2006 the hedge fund Amaranth lost \$6 billion because of trading risks taken by Brian Hunter. In 2008, Jérôme Kerviel lost over \$7 billion trading equity index futures for Société Générale. The huge losses at Daiwa, Shell, and Sumitomo were also each the result of the activities of a single individual.

The losses should not be viewed as an indictment of the whole derivatives industry. The derivatives market is a vast multitrillion dollar market that by most measures has been outstandingly successful and has served the needs of its users well. To quote from Alan Greenspan (May 2003):

The use of a growing array of derivatives and the related application of more sophisticated methods for measuring and managing risk are key factors underpinning the enhanced resilience of our largest financial intermediaries.

The events listed in Business Snapshots 35.1 and 35.2 represent a tiny proportion of the total trades (both in number and value). Nevertheless, it is worth considering carefully the lessons that can be learned from them.

35.1 LESSONS FOR ALL USERS OF DERIVATIVES

First, we consider the lessons appropriate to all users of derivatives, whether they are financial or nonfinancial companies.

Business Snapshot 35.1 Big Losses by Financial Institutions*Allied Irish bank*

This bank lost about \$700 million from speculative activities of one of its foreign exchange traders, John Rusnak, that lasted a number of years. Rusnak managed to cover up his losses by creating fictitious option trades.

Amaranth

This hedge fund lost \$6 billion in 2006 betting on the future direction of natural gas prices.

Barings

This 200-year-old British bank was destroyed in 1995 by the activities of one trader, Nick Leeson, in Singapore, who made big bets on the future direction of the Nikkei 225 using futures and options. The total loss was close to \$1 billion.

Daiwa Bank

A trader working in New York for this Japanese bank lost more than \$1 billion in the 1990s.

Enron's counterparties

Enron managed to conceal its true situation from its shareholders with some creative contracts. Several financial institutions that allegedly helped Enron do this have settled shareholder lawsuits for over \$1 billion.

Kidder Peabody (see page 106)

The activities of a single trader, Joseph Jett, led to this New York investment dealer losing \$350 million trading US government securities. The loss arose because of a mistake in the way the company's computer system calculated profits.

Long-Term Capital Management (see page 31)

This hedge fund lost about \$4 billion in 1998 as a result of Russia's default on its debt and the resultant flight to quality. The New York Federal Reserve organized an orderly liquidation of the fund by arranging for 14 banks to invest in the fund.

Midland Bank

This British bank lost \$500 million in the early 1990s largely because of a wrong bet on the direction of interest rates. It was later taken over by the Hong Kong and Shanghai bank.

Société Générale (see page 17)

Jérôme Kerviel lost over \$7 billion speculating on the future direction of equity indices in January 2008.

Subprime Mortgage Losses (see Chapter 8)

In 2007 investors lost confidence in the structured products created from US subprime mortgages. This led to a "credit crunch" and losses of tens of billions of dollars by financial institutions such as UBS, Merrill Lynch, and Citigroup.

Business Snapshot 35.2 Big Losses by Nonfinancial Organizations*Allied Lyons*

The treasury department of this drinks and food company lost \$150 million in 1991 selling call options on the US dollar–sterling exchange rate.

Gibson Greetings

The treasury department of this greeting card manufacturer lost about \$20 million in 1994 trading highly exotic interest rate derivatives contracts with Bankers Trust. They later sued Bankers Trust and settled out of court.

Hammersmith and Fulham (see page 173)

This British Local Authority lost about \$600 million on sterling interest rate swaps and options in 1988. All its contracts were later declared null and void by the British courts, much to the annoyance of the banks on the other side of the transactions.

Metallgesellschaft (see page 67)

This German company entered into long-term contracts to supply oil and gasoline and hedged them by rolling over short-term futures contracts. It lost \$1.3 billion when it was forced to discontinue this activity.

Orange County (see page 87)

The activities of the treasurer, Robert Citron, led to this California municipality losing about \$2 billion in 1994. The treasurer was using derivatives to speculate that interest rates would not rise.

Procter & Gamble (see page 745)

The treasury department of this large US company lost about \$90 million in 1994 trading highly exotic interest rate derivatives contracts with Bankers Trust. It later sued Bankers Trust and settled out of court.

Shell

A single employee working in the Japanese subsidiary of this company lost \$1 billion dollars in unauthorized trading of currency futures.

Sumitomo

A single trader working for this Japanese company lost about \$2 billion in the copper spot, futures, and options market in the 1990s.

Define Risk Limits

It is essential that all companies define in a clear and unambiguous way limits to the financial risks that can be taken. They should then set up procedures for ensuring that the limits are obeyed. Ideally, overall risk limits should be set at board level. These should then be converted to limits applicable to the individuals responsible for managing particular risks. Daily reports should indicate the gain or loss that will be experienced for particular movements in market variables. These should be checked against the actual gains and losses that are experienced to ensure that the valuation procedures underlying the reports are accurate.

It is particularly important that companies monitor risks carefully when derivatives are used. This is because, as we saw in Chapter 1, derivatives can be used for hedging,

speculation, and arbitrage. Without close monitoring, it is impossible to know whether a derivatives trader has switched from being a hedger to a speculator or switched from being an arbitrageur to being a speculator. Barings and Société Générale are classic examples of what can go wrong. Nick Leeson's mandate at Barings and Jérôme Kerviel's at Société Générale were to carry out low-risk arbitrage trades. They both switched from being arbitrageurs to taking huge bets on the future direction of stock indices. Systems at their banks were so inadequate that nobody knew the full extent of what they were doing.

The argument here is not that no risks should be taken. A treasurer working for a corporation, or a trader in a financial institution, or a fund manager should be allowed to take positions on the future direction of relevant market variables. But the sizes of the positions that can be taken should be limited and the systems in place should accurately report the risks being taken.

Take the Risk Limits Seriously

What happens if an individual exceeds risk limits and makes a profit? This is a tricky issue for senior management. It is tempting to ignore violations of risk limits when profits result. However, this is shortsighted. It leads to a culture where risk limits are not taken seriously, and it paves the way for a disaster. In some of the situations listed in Business Snapshots 35.1 and 35.2, the companies had become complacent about the risks they were taking because they had taken similar risks in previous years and made profits.

The classic example here is Orange County. Robert Citron's activities in 1991–93 had been very profitable for Orange County, and the municipality had come to rely on his trading for additional funding. People chose to ignore the risks he was taking because he had produced profits. Unfortunately, the losses made in 1994 far exceeded the profits from previous years.

The penalties for exceeding risk limits should be just as great when profits result as when losses result. Otherwise, traders who make losses are liable to keep increasing their bets in the hope that eventually a profit will result and all will be forgiven.

Do Not Assume You Can Outguess the Market

Some traders are quite possibly better than others. But no trader gets it right all the time. A trader who correctly predicts the direction in which market variables will move 60% of the time is doing well. If a trader has an outstanding track record (as Robert Citron did in the early 1990s), it is likely to be a result of luck rather than superior trading skill.

Suppose that a financial institution employs 16 traders and one of those traders makes profits in every quarter of a year. Should the trader receive a good bonus? Should the trader's risk limits be increased? The answer to the first question is that inevitably the trader will receive a good bonus. The answer to the second question should be no. The chance of making a profit in four consecutive quarters from random trading is 0.5^4 or 1 in 16. This means that just by chance one of the 16 traders will "get it right" every single quarter of the year. It should not be assumed that the trader's luck will continue and the trader's risk limits should not be increased.

Do Not Underestimate the Benefits of Diversification

When a trader appears good at predicting a particular market variable, there is a tendency to increase the trader's limits. We have just argued that this is a bad idea because it is quite likely that the trader has been lucky rather than clever. However, let us suppose that a fund is really convinced that the trader has special talents. How undiversified should it allow itself to become in order to take advantage of the trader's special skills? The answer is that the benefits from diversification are huge, and it is unlikely that any trader is so good that it is worth foregoing these benefits to speculate heavily on just one market variable.

An example will illustrate the point here. Suppose that there are 20 stocks, each of which have an expected return of 10% per annum and a standard deviation of returns of 30%. The correlation between the returns from any two of the stocks is 0.2. By dividing an investment equally among the 20 stocks, an investor has an expected return of 10% per annum and standard deviation of returns of 14.7%. Diversification enables the investor to reduce risks by over half. Another way of expressing this is that diversification enables an investor to double the expected return per unit of risk taken. The investor would have to be extremely good at stock picking to get a better risk–return tradeoff by investing in just one stock.

Carry out Scenario Analyses and Stress Tests

The calculation of risk measures such as VaR should always be accompanied by scenario analyses and stress testing to obtain an understanding of what can go wrong. These were mentioned in Chapter 21. They are very important. Human beings have an unfortunate tendency to anchor on one or two scenarios when evaluating decisions. In 1993 and 1994, for example, Procter & Gamble and Gibson Greetings may have been so convinced that interest rates would remain low that they ignored the possibility of a 100-basis-point increase in their decision making.

It is important to be creative in the way scenarios are generated and to use the judgment of experienced managers. One approach is to look at 10 or 20 years of data and choose the most extreme events as scenarios. Sometimes there is a shortage of data on a key variable. It is then sensible to choose a similar variable for which much more data is available and use historical daily percentage changes in that variable as a proxy for possible daily percentage changes in the key variable. For example, if there is little data on the prices of bonds issued by a particular country, historical data on prices of bonds issued by other similar countries can be used to develop possible scenarios.

35.2 LESSONS FOR FINANCIAL INSTITUTIONS

We now move on to consider lessons that are primarily relevant to financial institutions.

Monitor Traders Carefully

In trading rooms there is a tendency to regard high-performing traders as “untouchable” and to not subject their activities to the same scrutiny as other traders. Apparently

Joseph Jett, Kidder Peabody's star trader of Treasury instruments, was often "too busy" to answer questions and discuss his positions with the company's risk managers.

It is important that all traders—particularly those making high profits—be fully accountable. It is important for the financial institution to know whether the high profits are being made by taking unreasonably high risks. It is also important to check that the financial institution's computer systems and pricing models are correct and are not being manipulated in some way.

Separate the Front, Middle, and Back Office

The *front office* in a financial institution consists of the traders who are executing trades, taking positions, and so forth. The *middle office* consists of risk managers who are monitoring the risks being taken. The *back office* is where the record keeping and accounting takes place. Some of the worst derivatives disasters have occurred because these functions were not kept separate. Nick Leeson controlled both the front and back office for Barings in Singapore and was, as a result, able to conceal the disastrous nature of his trades from his superiors in London for some time. Jérôme Kerviel had worked in Société Générale's back office before becoming a trader and took advantage of his knowledge of its systems to hide his positions.

Do Not Blindly Trust Models

Some of the large losses incurred by financial institutions arose because of the models and computer systems being used. We discussed how Kidder Peabody was misled by its own systems on page 106.

If large profits are reported when relatively simple trading strategies are followed, there is a good chance that the models underlying the calculation of the profits are wrong. Similarly, if a financial institution appears to be particularly competitive on its quotes for a particular type of deal, there is a good chance that it is using a different model from other market participants, and it should analyze what is going on carefully. To the head of a trading room, getting too much business of a certain type can be just as worrisome as getting too little business of that type.

Be Conservative in Recognizing Inception Profits

When a financial institution sells a highly exotic instrument to a nonfinancial corporation, the valuation can be highly dependent on the underlying model. For example, instruments with long-dated embedded interest rate options can be highly dependent on the interest rate model used. In these circumstances, a phrase used to describe the daily marking to market of the deal is *marking to model*. This is because there are no market prices for similar deals that can be used as a benchmark.

Suppose that a financial institution manages to sell an instrument to a client for \$10 million more than it is worth—or at least \$10 million more than its model says it is worth. The \$10 million is known as an *inception profit*. When should it be recognized? There appears to be quite a variation in what different investment banks do. Some recognize the \$10 million immediately, whereas others are much more conservative and recognize it slowly over the life of the deal.

Recognizing inception profits immediately is very dangerous. It encourages traders to use aggressive models, take their bonuses, and leave before the model and the value of the deal come under close scrutiny. It is much better to recognize inception profits slowly, so that traders have the motivation to investigate the impact of several different models and several different sets of assumptions before committing themselves to a deal.

Do Not Sell Clients Inappropriate Products

It is tempting to sell corporate clients inappropriate products, particularly when they appear to have an appetite for the underlying risks. But this is shortsighted. The most dramatic example of this is the activities of Bankers Trust (BT) in the period leading up to the spring of 1994. Many of BT's clients were persuaded to buy high-risk and totally inappropriate products. A typical product (e.g., the 5/30 swap discussed on page 745) would give the client a good chance of saving a few basis points on its borrowings and a small chance of costing a large amount of money. The products worked well for BT's clients in 1992 and 1993, but blew up in 1994 when interest rates rose sharply. The bad publicity that followed hurt BT greatly. The years it had spent building up trust among corporate clients and developing an enviable reputation for innovation in derivatives were largely lost as a result of the activities of a few overly aggressive salesmen. BT was forced to pay large amounts of money to its clients to settle lawsuits out of court. It was taken over by Deutsche Bank in 1999.

Do Not Ignore Liquidity Risk

Financial engineers usually base the pricing of exotic instruments and other instruments that trade relatively infrequently on the prices of actively traded instruments. For example:

1. A financial engineer often calculates a zero curve from actively traded government bonds (known as on-the-run bonds) and uses it to price bonds that trade less frequently (off-the-run bonds).
2. A financial engineer often implies the volatility of an asset from actively traded options and uses it to price less actively traded options.
3. A financial engineer often implies information about the behavior of interest rates from actively traded interest rate caps and swap options and uses it to price products that are highly structured.

These practices are not unreasonable. However, it is dangerous to assume that less actively traded instruments can always be traded at close to their theoretical price. When financial markets experience a shock of one sort or another there is often a "flight to quality." Liquidity becomes very important to investors, and illiquid instruments often sell at a big discount to their theoretical values. This happened in 2007 following the jolt to credit markets caused by lack of confidence in securities backed by subprime mortgages.

Another example of losses arising from liquidity risk is provided by Long-Term Capital Management (LTCM), which was discussed in Business Snapshot 2.2. This hedge fund followed a strategy known as *convergence arbitrage*. It attempted to identify two securities (or portfolios of securities) that should in theory sell for the same price. If the market price of one security was less than that of the other, it would buy that security

and sell the other. The strategy is based on the idea that if two securities have the same theoretical price their market prices should eventually be the same.

In the summer of 1998 LTCM made a huge loss. This was largely because a default by Russia on its debt caused a flight to quality. LTCM tended to be long illiquid instruments and short the corresponding liquid instruments (for example, it was long off-the-run bonds and short on-the-run bonds). The spreads between the prices of illiquid instruments and the corresponding liquid instruments widened sharply after the Russian default. LTCM was highly leveraged. It experienced huge losses and there were margin calls on its positions that it found difficult to meet.

The LTCM story reinforces the importance of carrying out scenario analyses and stress testing to look at what can happen in the worst of all worlds. LTCM could have tried to examine other times in history when there have been extreme flights to quality to quantify the liquidity risks it was facing.

Beware When Everyone Is Following the Same Trading Strategy

It sometimes happens that many market participants are following essentially the same trading strategy. This creates a dangerous environment where there are liable to be big market moves, unstable markets, and large losses for the market participants.

We gave one example of this in Chapter 18 when discussing portfolio insurance and the market crash of October 1987. In the months leading up to the crash, increasing numbers of portfolio managers were attempting to insure their portfolios by creating synthetic put options. They bought stocks or stock index futures after a rise in the market and sold them after a fall. This created an unstable market. A relatively small decline in stock prices could lead to a wave of selling by portfolio insurers. The latter would lead to a further decline in the market, which could give rise to another wave of selling, and so on. There is little doubt that without portfolio insurance the crash of October 1987 would have been much less severe.

Another example is provided by LTCM in 1998. Its position was made more difficult by the fact that many other hedge funds were following similar convergence arbitrage strategies. After the Russian default and the flight to quality, LTCM tried to liquidate part of its portfolio to meet margin calls. Unfortunately, other hedge funds were facing similar problems to LTCM and trying to do similar trades. This exacerbated the situation, causing liquidity spreads to be even higher than they would otherwise have been and reinforcing the flight to quality. Consider, for example, LTCM's position in US Treasury bonds. It was long the illiquid off-the-run bonds and short the liquid on-the-run bonds. When a flight to quality caused spreads between yields on the two types of bonds to widen, LTCM had to liquidate its positions by selling off-the-run bonds and buying on-the-run bonds. Other large hedge funds were doing the same. As a result, the price of on-the-run bonds rose relative to off-the-run bonds and the spread between the two yields widened even more than it had done already.

A further example is provided by the activities of British insurance companies in the late 1990s. These insurance companies had entered into many contracts promising that the rate of interest applicable to an annuity received by an individual on retirement would be the greater of the market rate and a guaranteed rate. At about the same time, all insurance companies decided to hedge part of their risks on these contracts by buying long-dated swap options from financial institutions. The financial institutions they dealt with hedged their risks by buying huge numbers of long-dated sterling

bonds. As a result, bond prices rose and long sterling rates declined. More bonds had to be bought to maintain the dynamic hedge, long sterling rates declined further, and so on. Financial institutions lost money and, because long rates declined, insurance companies found themselves in a worse position on the risks that they had chosen not to hedge.

The chief lesson to be learned from these stories is that it is important to see the big picture of what is going on in financial markets and to understand the risks inherent in situations where many market participants are following the same trading strategy.

Short-Term Funding Can Create Liquidity Problems

The interest rate risks when a bank funds long-term assets with short-term liabilities are now well understood (see Section 4.10). Banks are usually careful to hedge these risks with interest rate swaps or other derivatives.

The liquidity risks when a financial institution funds long-term assets (or other long-term needs) with short-term liabilities are in many ways more serious than the interest rate risks, but received less attention until the onset of the credit crisis in 2007. The problem is that, when the market (rightly or wrongly) loses confidence in a financial institution, the financial institution will find it impossible to roll over its liabilities. In the normal course of events, a financial institution might issue 1-month commercial paper on July 1, repay it on August 1 with a new issue of 1-month commercial paper, repay the new issue on September 1 with yet another issue of 1-month commercial paper, and so on. When there is a loss of confidence, new commercial paper cannot be issued and there is an immediate liquidity problem.

The credit crisis created a loss of confidence in many financial institutions, particularly those heavily involved in mortgage lending or those thought to have big positions in the tranches created from subprime mortgages. Northern Rock, a mortgage lender in the UK, was one of the first casualties of the crisis (see Business Snapshot 4.3). It financed much of its mortgage lending with short-term commercial paper. When investors lost confidence in the real-estate market, the commercial paper could not be rolled over. As mentioned in Business Snapshot 1.1, Lehman also financed much of its long-term needs with short-term paper. When there was concern about its health, the short-term paper could not be rolled over, accelerating the company's bankruptcy. The companies that were in the business of creating the products discussed in Chapter 8 from subprime mortgages also experienced liquidity problems because their (long-term) needs to finance their inventories of mortgages that were awaiting securitization were usually financed with short-term paper.

One of the results of the credit crisis is that a bank's supervisors now monitor its liquidity as well as its capital adequacy.

Market Transparency Is Important

One of the lessons from the credit crunch of 2007 is that market transparency is important. During the period leading up to 2007, investors traded highly structured products without any real knowledge of the underlying assets. All they knew was the credit rating of the security being traded. With hindsight, we can say that investors should have demanded more information about the underlying assets and should have more carefully assessed the risks they were taking—but it is easy to be wise after the event!

The subprime meltdown of August 2007 caused investors to lose confidence in all structured products and withdraw from that market. This led to a market breakdown where tranches of structured products could only be sold at prices well below their theoretical values. There was a flight to quality and credit spreads increased. If there had been market transparency so that investors understood the asset-backed securities they were buying, there would still have been subprime losses, but the flight to quality and disruptions to the market would have been less pronounced.

Manage Incentives

A key lesson from the credit crisis of 2007 and 2008 is the importance of incentives. The bonus systems in banks tend to emphasize short-term performance. Some financial institutions have switched to systems where bonuses are based on performance over a longer window than one year (for example, five years). This has obvious advantages. It discourages traders from doing trades that will look good in the short run, but may “blow up” in a few years.

When loans are securitized, it is important to align the interests of the party originating the loan with the party who bears the ultimate risk so that the originator does not have an incentive to misrepresent the loan. One way of doing this is for regulators to require the originator of a loan portfolio to keep a stake in all the tranches and other instruments that are created from the portfolio.

Never Ignore Risk Management

When times are good (or appear to be good), there is a tendency to assume that nothing can go wrong and ignore the output from stress tests and other analyses carried out by the risk management group. There are many stories of risk managers not being listened to in the period leading up to the credit crisis of 2007. The comment of Chuck Prince, CEO of Citigroup, in July 2007 (just before the credit crisis) provides an example of exactly the wrong attitude to risk management:

When the music stops, in terms of liquidity, things will be complicated. But as long as the music is playing, you've got to get up and dance. We're still dancing.

Mr. Prince lost his job later in the year and Citigroup's losses from the credit crisis were over \$50 billion.

35.3 LESSONS FOR NONFINANCIAL CORPORATIONS

We now consider lessons primarily applicable to nonfinancial corporations.

Make Sure You Fully Understand the Trades You Are Doing

Corporations should never undertake a trade or a trading strategy that they do not fully understand. This is a somewhat obvious point, but it is surprising how often a trader working for a nonfinancial corporation will, after a big loss, admit to not knowing what was really going on and claim to have been misled by investment bankers. Robert Citron, the treasurer of Orange County did this. So did the traders working for Hammersmith and Fulham, who in spite of their huge positions were surprisingly uninformed about how the swaps and other interest rate derivatives they traded really worked.

If a senior manager in a corporation does not understand a trade proposed by a subordinate, the trade should not be approved. A simple rule of thumb is that if a trade and the rationale for entering into it are so complicated that they cannot be understood by the manager, it is almost certainly inappropriate for the corporation. The trades undertaken by Procter & Gamble and Gibson Greetings would have been vetoed using this criterion.

One way of ensuring that you fully understand a financial instrument is to value it. If a corporation does not have the in-house capability to value an instrument, it should not trade it. In practice, corporations often rely on their derivatives dealers for valuation advice. This is dangerous, as Procter & Gamble and Gibson Greetings found out. When they wanted to unwind their deals, they found they were facing prices produced by Bankers Trust's proprietary models, which they had no way of checking.

Make Sure a Hedger Does Not Become a Speculator

One of the unfortunate facts of life is that hedging is relatively dull, whereas speculation is exciting. When a company hires a trader to manage foreign exchange, commodity price, or interest rate risk, there is a danger that the following might happen. At first, the trader does the job diligently and earns the confidence of top management. He or she assesses the company's exposures and hedges them. As time goes by, the trader becomes convinced that he or she can outguess the market. Slowly the trader becomes a speculator. At first things go well, but then a loss is made. To recover the loss, the trader doubles up the bets. Further losses are made—and so on. The result is likely to be a disaster.

As mentioned earlier, clear limits to the risks that can be taken should be set by senior management. Controls should be put in place to ensure that the limits are obeyed. The trading strategy for a corporation should start with an analysis of the risks facing the corporation in foreign exchange, interest rate, commodity markets, and so on. A decision should then be taken on how the risks are to be reduced to acceptable levels. It is a clear sign that something is wrong within a corporation if the trading strategy is not derived in a very direct way from the company's exposures.

Be Cautious about Making the Treasury Department a Profit Center

In the last 20 years there has been a tendency to make the treasury department within a corporation a profit center. This appears to have much to recommend it. The treasurer is motivated to reduce financing costs and manage risks as profitably as possible. The problem is that the potential for the treasurer to make profits is limited. When raising funds and investing surplus cash, the treasurer is facing an efficient market. The treasurer can usually improve the bottom line only by taking additional risks. The company's hedging program gives the treasurer some scope for making shrewd decisions that increase profits. But it should be remembered that the goal of a hedging program is to reduce risks, not to increase expected profits. As pointed out in Chapter 3, the decision to hedge will lead to a worse outcome than the decision not to hedge roughly 50% of the time. The danger of making the treasury department a profit center is that the treasurer is motivated to become a speculator. This is liable to lead to the type of outcome experienced by Orange County, Procter & Gamble, or Gibson Greetings.

SUMMARY

The huge losses experienced from the use of derivatives have made many treasurers very wary. Following some of the losses, some nonfinancial corporations have announced plans to reduce or even eliminate their use of derivatives. This is unfortunate because derivatives provide treasurers with very efficient ways to manage risks.

The stories behind the losses emphasize the point, made as early as Chapter 1, that derivatives can be used for either hedging or speculation; that is, they can be used either to reduce risks or to take risks. Most losses occurred because derivatives were used inappropriately. Employees who had an implicit or explicit mandate to hedge their company's risks decided instead to speculate.

The key lesson to be learned from the losses is the importance of *internal controls*. Senior management within a company should issue a clear and unambiguous policy statement about how derivatives are to be used and the extent to which it is permissible for employees to take positions on movements in market variables. Management should then institute controls to ensure that the policy is carried out. It is a recipe for disaster to give individuals authority to trade derivatives without a close monitoring of the risks being taken.

FURTHER READING

- Dunbar, N. *Inventing Money: The Story of Long-Term Capital Management and the Legends Behind It*. Chichester, UK: Wiley, 2000.
- Jorion, P. *Big Bets Gone Bad: Derivatives and Bankruptcy in Orange County*. New York: Academic Press, 1995.
- Jorion, P. "How Long-Term Lost Its Capital," *Risk* (September 1999).
- Ju, X., and N. Pearson. "Using Value at Risk to Control Risk Taking: How Wrong Can You Be?" *Journal of Risk*, 1 (1999): 5–36.
- Persaud, A. D. (ed.) *Liquidity Black Holes: Understanding, Quantifying and Managing Financial Liquidity Risk*. London, Risk Books, 2003.
- Sorkin, A. R., *Too Big to Fail*. New York: Penguin, 2009.
- Tett, G. *Fool's Gold: How the Bold Dream of a Small Tribe at J. P. Morgan Was Corrupted by Wall Street Greed and Unleashed a Catastrophe*. New York: Free Press, 2009.
- Thomson, R. *Apocalypse Roulette: The Lethal World of Derivatives*. London: Macmillan, 1998.
- Zhang, P. G. *Barings, Bankruptcy and Financial Derivatives*. Singapore: World Scientific Publishing, 1995.



Glossary of Terms

ABS *See* Asset-Backed Security.

ABS CDO Instrument where tranches are created from the tranches of ABSs.

Accrual Swap An interest rate swap where interest on one side accrues only when a certain condition is met.

Accrued Interest The interest earned on a bond since the last coupon payment date.

Adaptive Mesh Model A model developed by Figlewski and Gao that grafts a high-resolution tree on to a low-resolution tree so that there is more detailed modeling of the asset price in critical regions.

Agency Costs Costs arising from a situation where the agent (e.g., manager) is not motivated to act in the best interests of the principal (e.g., shareholder).

American Option An option that can be exercised at any time during its life.

Amortizing Swap A swap where the notional principal decreases in a predetermined way as time passes.

Analytic Result Result where answer is in the form of an equation.

Arbitrage A trading strategy that takes advantage of two or more securities being mispriced relative to each other.

Arbitrageur An individual engaging in arbitrage.

Asian Option An option with a payoff dependent on the average price of the underlying asset during a specified period.

Ask Price The price that a dealer is offering to sell an asset.

Asked Price *See* Ask Price.

Asset-Backed Security Security created from a portfolio of loans, bonds, credit card receivables, or other assets.

Asset-or-Nothing Call Option An option that provides a payoff equal to the asset price if the asset price is above the strike price and zero otherwise.

Asset-or-Nothing Put Option An option that provides a payoff equal to the asset price if the asset price is below the strike price and zero otherwise.

Asset Swap Exchanges the coupon on a bond for LIBOR plus a spread.

- As-You-Like-It Option** *See* Chooser Option.
- At-the-Money Option** An option in which the strike price equals the price of the underlying asset.
- Average Price Call Option** An option giving a payoff equal to the greater of zero and the amount by which the average price of the asset exceeds the strike price.
- Average Price Put Option** An option giving a payoff equal to the greater of zero and the amount by which the strike price exceeds the average price of the asset.
- Average Strike Option** An option that provides a payoff dependent on the difference between the final asset price and the average asset price.
- Backdating** Practice (often illegal) of marking a document with a date that precedes the current date.
- Back Testing** Testing a value-at-risk or other model using historical data.
- Backwards Induction** A procedure for working from the end of a tree to its beginning in order to value an option.
- Barrier Option** An option whose payoff depends on whether the path of the underlying asset has reached a barrier (i.e., a certain predetermined level).
- Base Correlation** Correlation that leads to the price of a 0% to X% CDO tranche being consistent with the market for a particular value of X.
- Basel Committee** Committee responsible for regulation of banks internationally.
- Basis** The difference between the spot price and the futures price of a commodity.
- Basis Point** When used to describe an interest rate, a basis point is one hundredth of one percent (= 0.01%)
- Basis Risk** The risk to a hedger arising from uncertainty about the basis at a future time.
- Basis Swap** A swap where cash flows determined by one floating reference rate are exchanged for cash flows determined by another floating reference rate.
- Basket Credit Default Swap** Credit default swap where there are several reference entities.
- Basket Option** An option that provides a payoff dependent on the value of a portfolio of assets.
- Bear Spread** A short position in a put option with strike price K_1 combined with a long position in a put option with strike price K_2 where $K_2 > K_1$. (A bear spread can also be created with call options.)
- Bermudan Option** An option that can be exercised on specified dates during its life.
- Beta** A measure of the systematic risk of an asset.
- Bid–Ask Spread** The amount by which the ask price exceeds the bid price.
- Bid–Offer Spread** *See* Bid–Ask Spread.
- Bid Price** The price that a dealer is prepared to pay for an asset.
- Binary Credit Default Swap** Instrument where there is a fixed dollar payoff in the event of a default by a particular company.
- Binary Option** Option with a discontinuous payoff, e.g., a cash-or-nothing option or an asset-or-nothing option.

- Binomial Model** A model where the price of an asset is monitored over successive short periods of time. In each short period it is assumed that only two price movements are possible.
- Binomial Tree** A tree that represents how an asset price can evolve under the binomial model.
- Bivariate Normal Distribution** A distribution for two correlated variables, each of which is normal.
- Black's Approximation** An approximate procedure developed by Fischer Black for valuing a call option on a dividend-paying stock.
- Black's Model** An extension of the Black–Scholes model for valuing European options on futures contracts. As described in Chapter 26, it is used extensively in practice to value European options when the distribution of the asset price at maturity is assumed to be lognormal.
- Black–Scholes–Merton Model** A model for pricing European options on stocks, developed by Fischer Black, Myron Scholes, and Robert Merton.
- Bond Option** An option where a bond is the underlying asset.
- Bond Yield** Discount rate which, when applied to all the cash flows of a bond, causes the present value of the cash flows to equal the bond's market price.
- Bootstrap Method** A procedure for calculating the zero-coupon yield curve from market data.
- Boston Option** *See* Deferred Payment Option.
- Box Spread** A combination of a bull spread created from calls and a bear spread created from puts.
- Break Forward** *See* Deferred Payment Option.
- Brownian Motion** *See* Wiener Process.
- Bull Spread** A long position in a call with strike price K_1 combined with a short position in a call with strike price K_2 , where $K_2 > K_1$. (A bull spread can also be created with put options.)
- Butterfly Spread** A position that is created by taking a long position in a call with strike price K_1 , a long position in a call with strike price K_3 , and a short position in two calls with strike price K_2 , where $K_3 > K_2 > K_1$ and $K_2 = 0.5(K_1 + K_3)$. (A butterfly spread can also be created with put options.)
- Calendar Spread** A position that is created by taking a long position in a call option that matures at one time and a short position in a similar call option that matures at a different time. (A calendar spread can also be created using put options.)
- Calibration** Method for implying a model's parameters from the prices of actively traded options.
- Callable Bond** A bond containing provisions that allow the issuer to buy it back at a predetermined price at certain times during its life.
- Call Option** An option to buy an asset at a certain price by a certain date.
- Cancelable Swap** Swap that can be canceled by one side on prespecified dates.
- Cap** *See* Interest Rate Cap.

- Cap Rate** The rate determining payoffs in an interest rate cap.
- Capital Asset Pricing Model** A model relating the expected return on an asset to its beta.
- Caplet** One component of an interest rate cap.
- Case-Shiller Index** Index of house prices in the United States.
- Cash Flow Mapping** A procedure for representing an instrument as a portfolio of zero-coupon bonds for the purpose of calculating value at risk.
- Cash-or-Nothing Call Option** An option that provides a fixed predetermined payoff if the final asset price is above the strike price and zero otherwise.
- Cash-or-Nothing Put Option** An option that provides a fixed predetermined payoff if the final asset price is below the strike price and zero otherwise.
- Cash Settlement** Procedure for settling a futures contract in cash rather than by delivering the underlying asset.
- CAT Bond** Bond where the interest and, possibly, the principal paid are reduced if a particular category of “catastrophic” insurance claims exceed a certain amount.
- CDD** Cooling degree days. The maximum of zero and the amount by which the daily average temperature is greater than 65° Fahrenheit. The average temperature is the average of the highest and lowest temperatures (midnight to midnight).
- CDO** *See* Collateralized Debt Obligation.
- CDO Squared** An instrument in which the default risks in a portfolio of CDO tranches are allocated to new securities.
- CDX NA IG** Portfolio of 125 North American companies.
- Central Clearing Party** A clearing house used for over-the-counter contracts.
- Cheapest-to-Deliver Bond** The bond that is cheapest to deliver in the Chicago Board of Trade bond futures contract.
- Cholesky Decomposition** A method of sampling from a multivariate normal distribution.
- Chooser Option** An option where the holder has the right to choose whether it is a call or a put at some point during its life.
- Class of Options** *See* Option Class.
- Clean Price of Bond** The quoted price of a bond. The cash price paid for the bond (or dirty price) is calculated by adding the accrued interest to the clean price.
- Clearing House** A firm that guarantees the performance of the parties in a derivatives transaction (also referred to as a clearing corporation).
- Clearing Margin** A margin posted by a member of a clearinghouse.
- Cliquet Option** A series of call or put options with rules for determining strike prices. Typically, one option starts when the previous one terminates.
- CMO** Collateralized Mortgage Obligation.
- Collar** *See* Interest Rate Collar.
- Collateralization** A system for posting collateral by one or both parties in a derivatives transaction.

- Collateralized Debt Obligation** A way of packaging credit risk. Several classes of securities (known as tranches) are created from a portfolio of bonds and there are rules for determining how the cost of defaults are allocated to classes.
- Collateralized Mortgage Obligation (CMO)** A mortgage-backed security where investors are divided into classes and there are rules for determining how principal repayments are channeled to the classes.
- Combination** A position involving both calls and puts on the same underlying asset.
- Commodity Futures Trading Commission** A body that regulates trading in futures contracts in the United States.
- Commodity Swap** A swap where cash flows depend on the price of a commodity.
- Compound Correlation** Correlation implied from the market price of a CDO tranche.
- Compound Option** An option on an option.
- Compounding Frequency** This defines how an interest rate is measured.
- Compounding Swap** Swap where interest compounds instead of being paid.
- Conditional Value at Risk (C-VaR)** Expected loss during N days conditional on being in the $(100 - X)\%$ tail of the distribution of profits/losses. The variable N is the time horizon and $X\%$ is the confidence level.
- Confirmation** Contract confirming verbal agreement between two parties to a trade in the over-the-counter market.
- Constant Elasticity of Variance (CEV) Model** Model where the variance of the change in a variable in a short period of time is proportional to the value of the variable.
- Constant Maturity Swap (CMS)** A swap where a swap rate is exchanged for either a fixed rate or a floating rate on each payment date.
- Constant Maturity Treasury Swap** A swap where the yield on a Treasury bond is exchanged for either a fixed rate or a floating rate on each payment date.
- Consumption Asset** An asset held for consumption rather than investment.
- Contango** A situation where the futures price is above the expected future spot price (also often used to refer to the situation where the futures price is above the current spot price).
- Continuous Compounding** A way of quoting interest rates. It is the limit as the assumed compounding interval is made smaller and smaller.
- Control Variate Technique** A technique that can sometimes be used for improving the accuracy of a numerical procedure.
- Convenience Yield** A measure of the benefits from ownership of an asset that are not obtained by the holder of a long futures contract on the asset.
- Conversion Factor** A factor used to determine the number of bonds that must be delivered in the Chicago Board of Trade bond futures contract.
- Convertible Bond** A corporate bond that can be converted into a predetermined amount of the company's equity at certain times during its life.
- Convexity** A measure of the curvature in the relationship between bond prices and bond yields.

- Convexity Adjustment** An overworked term. For example, it can refer to the adjustment necessary to convert a futures interest rate to a forward interest rate. It can also refer to the adjustment to a forward rate that is sometimes necessary when Black's model is used.
- Copula** A way of defining the correlation between variables with known distributions.
- Cornish–Fisher Expansion** An approximate relationship between the fractiles of a probability distribution and its moments.
- Cost of Carry** The storage costs plus the cost of financing an asset minus the income earned on the asset.
- Counterparty** The opposite side in a financial transaction.
- Coupon** Interest payment made on a bond.
- Covariance** Measure of the linear relationship between two variables (equals the correlation between the variables times the product of their standard deviations).
- Covariance Matrix** *See* Variance–Covariance Matrix.
- Covered Call** A short position in a call option on an asset combined with a long position in the asset.
- Crashophobia** Fear of a stock market crash that some people claim causes the market to increase the price of deep-out-of-the-money put options.
- Credit Contagion** The tendency of a default by one company to lead to defaults by other companies.
- Credit Default Swap** An instrument that gives the holder the right to sell a bond for its face value in the event of a default by the issuer.
- Credit Derivative** A derivative whose payoff depends on the creditworthiness of one or more companies or countries.
- Credit Index** Index that tracks the cost of buying protection for each company in a portfolio (e.g., CDX NA IG and iTraxx Europe).
- Credit Rating** A measure of the creditworthiness of a bond issue.
- Credit Ratings Transition Matrix** A table showing the probability that a company will move from one credit rating to another during a certain period of time.
- Credit Risk** The risk that a loss will be experienced because of a default by the counterparty in a derivatives transaction.
- Credit Spread Option** Option whose payoff depends on the spread between the yields earned on two assets.
- Credit Value Adjustment** Adjustment to value of derivatives outstanding with a counterparty to reflect the counterparty's default risk.
- Credit Value at Risk** The credit loss that will not be exceeded at some specified confidence level.
- CreditMetrics** A procedure for calculating credit value at risk.
- Cross Hedging** Hedging an exposure to the price of one asset with a contract on another asset.
- Cumulative Distribution Function** The probability that a variable will be less than x as a function of x .

- Currency Swap** A swap where interest and principal in one currency are exchanged for interest and principal in another currency.
- CVA** *See* Credit Value Adjustment.
- Day Count** A convention for quoting interest rates.
- Day Trade** A trade that is entered into and closed out on the same day.
- Default Correlation** Measures the tendency of two companies to default at about the same time.
- Default Intensity** *See* Hazard Rate.
- Default Probability Density** Measures the unconditional probability of default in a future short period of time.
- Deferred Payment Option** An option where the price paid is deferred until the end of the option's life.
- Deferred Swap** An agreement to enter into a swap at some time in the future (also called a forward swap).
- Delivery Price** Price agreed to (possibly some time in the past) in a forward contract.
- Delta** The rate of change of the price of a derivative with the price of the underlying asset.
- Delta Hedging** A hedging scheme that is designed to make the price of a portfolio of derivatives insensitive to small changes in the price of the underlying asset.
- Delta-Neutral Portfolio** A portfolio with a delta of zero so that there is no sensitivity to small changes in the price of the underlying asset.
- DerivaGem** The software accompanying this book.
- Derivative** An instrument whose price depends on, or is derived from, the price of another asset.
- Deterministic Variable** A variable whose future value is known.
- Diagonal Spread** A position in two calls where both the strike prices and times to maturity are different. (A diagonal spread can also be created with put options.)
- Differential Swap** A swap where a floating rate in one currency is exchanged for a floating rate in another currency and both rates are applied to the same principal.
- Diffusion Process** Model where value of asset changes continuously (no jumps).
- Dirty Price of Bond** Cash price of bond.
- Discount Bond** *See* Zero-Coupon Bond.
- Discount Instrument** An instrument, such as a Treasury bill, that provides no coupons.
- Diversification** Reducing risk by dividing a portfolio between many different assets.
- Dividend** A cash payment made to the owner of a stock.
- Dividend Yield** The dividend as a percentage of the stock price.
- Dollar Duration** The product of a bond's modified duration and the bond price.
- Down-and-In Option** An option that comes into existence when the price of the underlying asset declines to a prespecified level.

- Down-and-Out Option** An option that ceases to exist when the price of the underlying asset declines to a prespecified level.
- Downgrade Trigger** A clause in a contract that states that the contract will be terminated with a cash settlement if the credit rating of one side falls below a certain level.
- Drift Rate** The average increase per unit of time in a stochastic variable.
- Duration** A measure of the average life a bond. It is also an approximation to the ratio of the proportional change in the bond price to the absolute change in its yield.
- Duration Matching** A procedure for matching the durations of assets and liabilities in a financial institution.
- DV01** The dollar value of a 1-basis-point increase in all interest rates.
- Dynamic Hedging** A procedure for hedging an option position by periodically changing the position held in the underlying asset. The objective is usually to maintain a delta-neutral position.
- Early Exercise** Exercise prior to the maturity date.
- Efficient Market Hypothesis** A hypothesis that asset prices reflect relevant information.
- Electronic Trading** System of trading where a computer is used to match buyers and sellers.
- Embedded Option** An option that is an inseparable part of another instrument.
- Empirical Research** Research based on historical market data.
- Employee Stock Option** A stock option issued by company on its own stock and given to its employees as part of their remuneration.
- Equilibrium Model** A model for the behavior of interest rates derived from a model of the economy.
- Equity Swap** A swap where the return on an equity portfolio is exchanged for either a fixed or a floating rate of interest.
- Equity Tranche** The tranche that first absorbs losses.
- Eurocurrency** A currency that is outside the formal control of the issuing country's monetary authorities.
- Eurodollar** A dollar held in a bank outside the United States.
- Eurodollar Futures Contract** A futures contract written on a Eurodollar deposit.
- Eurodollar Interest Rate** The interest rate on a Eurodollar deposit.
- European Option** An option that can be exercised only at the end of its life.
- EWMA** Exponentially weighted moving average.
- Exchange Option** An option to exchange one asset for another.
- Ex-dividend Date** When a dividend is declared, an ex-dividend date is specified. Investors who own shares of the stock just before the ex-dividend date receive the dividend.
- Exercise Limit** Maximum number of option contracts that can be exercised within a five-day period.

- Exercise Multiple** Ratio of stock price to strike price at time of exercise for employee stock option.
- Exercise Price** The price at which the underlying asset may be bought or sold in an option contract (also called the strike price).
- Exotic Option** A nonstandard option.
- Expectations Theory** The theory that forward interest rates equal expected future spot interest rates.
- Expected Shortfall** *See* Conditional Value at Risk.
- Expected Value of a Variable** The average value of the variable obtained by weighting the alternative values by their probabilities.
- Expiration Date** The end of life of a contract.
- Explicit Finite Difference Method** A method for valuing a derivative by solving the underlying differential equation. The value of the derivative at time t is related to three values at time $t + \Delta t$. It is essentially the same as the trinomial tree method.
- Exponentially Weighted Moving Average Model** A model where exponential weighting is used to provide forecasts for a variable from historical data. It is sometimes applied to variances and covariances in value at risk calculations.
- Exponential Weighting** A weighting scheme where the weight given to an observation depends on how recent it is. The weight given to an observation i time periods ago is λ times the weight given to an observation $i - 1$ time periods ago where $\lambda < 1$.
- Exposure** The maximum loss from default by a counterparty.
- Extendable Bond** A bond whose life can be extended at the option of the holder.
- Extendable Swap** A swap whose life can be extended at the option of one side to the contract.
- Factor** Source of uncertainty.
- Factor analysis** An analysis aimed at finding a small number of factors that describe most of the variation in a large number of correlated variables (similar to a principal components analysis).
- FAS 123** Accounting standard in United States relating to employee stock options.
- FAS 133** Accounting standard in United States relating to instruments used for hedging.
- FASB** Financial Accounting Standards Board.
- FICO** A credit score developed by Fair Isaac Corporation.
- Financial Intermediary** A bank or other financial institution that facilitates the flow of funds between different entities in the economy.
- Finite Difference Method** A method for solving a differential equation.
- Flat Volatility** The name given to volatility used to price a cap when the same volatility is used for each caplet.
- Flex Option** An option traded on an exchange with terms that are different from the standard options traded by the exchange.
- Flexi Cap** Interest rate cap where there is a limit on the total number of caplets that can be exercised.

- Floor** *See* Interest Rate Floor.
- Floor–Ceiling Agreement** *See* Collar.
- Floorlet** One component of a floor.
- Floor Rate** The rate in an interest rate floor agreement.
- Foreign Currency Option** An option on a foreign exchange rate.
- Forward Contract** A contract that obligates the holder to buy or sell an asset for a predetermined delivery price at a predetermined future time.
- Forward Exchange Rate** The forward price of one unit of a foreign currency.
- Forward Interest Rate** The interest rate for a future period of time implied by the rates prevailing in the market today.
- Forward Price** The delivery price in a forward contract that causes the contract to be worth zero.
- Forward Rate** Rate of interest for a period of time in the future implied by today's zero rates.
- Forward Rate Agreement (FRA)** Agreement that a certain interest rate will apply to a certain principal amount for a certain time period in the future.
- Forward Risk-Neutral World** A world is forward risk-neutral with respect to a certain asset when the market price of risk equals the volatility of that asset.
- Forward Start Option** An option designed so that it will be at-the-money at some time in the future.
- Forward Swap** *See* Deferred Swap.
- Futures Commission Merchants** Futures traders who are following instructions from clients.
- Futures Contract** A contract that obligates the holder to buy or sell an asset at a predetermined delivery price during a specified future time period. The contract is settled daily.
- Futures Option** An option on a futures contract.
- Futures Price** The delivery price currently applicable to a futures contract.
- Futures-Style Option** Futures contract on the payoff from an option.
- Gamma** The rate of change of delta with respect to the asset price.
- Gamma-Neutral Portfolio** A portfolio with a gamma of zero.
- Gap Option** European call or put option where there are two strike prices. One determines whether the option is exercised. The other determines the payoff.
- GARCH Model** A model for forecasting volatility where the variance rate follows a mean-reverting process.
- Gaussian Copula Model** A model for defining a correlation structure between two or more variables. In some credit derivatives models, it is used to define a correlation structure for times to default.
- Gaussian Quadrature** Procedure for integrating over a normal distribution.
- Generalized Wiener Process** A stochastic process where the change in a variable in time t has a normal distribution with mean and variance both proportional to t .

- Geometric Average** The n th root of the product of n numbers.
- Geometric Brownian Motion** A stochastic process often assumed for asset prices where the logarithm of the underlying variable follows a generalized Wiener process.
- Girsanov's Theorem** Result showing that when we change the measure (e.g., move from real world to risk-neutral world) the expected return of a variable changes but the volatility remains the same.
- Greeks** Hedge parameters such as delta, gamma, vega, theta, and rho.
- Haircut** Discount applied to the value of an asset for collateral purposes.
- Hazard Rate** Measures probability of default in a short period of time conditional on no earlier default.
- HDD** Heating degree days. The maximum of zero and the amount by which the daily average temperature is less than 65° Fahrenheit. The average temperature is the average of the highest and lowest temperatures (midnight to midnight).
- Hedge** A trade designed to reduce risk.
- Hedge Funds** Funds that are subject to less regulation and fewer restrictions than mutual funds. They can take short positions and use derivatives, but they cannot publicly offer their securities.
- Hedger** An individual who enters into hedging trades.
- Hedge Ratio** The ratio of the size of a position in a hedging instrument to the size of the position being hedged.
- Historical Simulation** A simulation based on historical data.
- Historic Volatility** A volatility estimated from historical data.
- Holiday Calendar** Calendar defining which days are holidays for the purposes of determining payment dates in a swap.
- IMM Dates** Third Wednesday in March, June, September, and December.
- Implicit Finite Difference Method** A method for valuing a derivative by solving the underlying differential equation. The value of the derivative at time $t + \Delta t$ is related to three values at time t .
- Implied Correlation** Correlation number implied from the price of a credit derivative using the Gaussian copula or similar model.
- Implied Distribution** A distribution for a future asset price implied from option prices.
- Implied Tree** A tree describing the movements of an asset price that is constructed to be consistent with observed option prices.
- Implied Volatility** Volatility implied from an option price using the Black–Scholes or a similar model.
- Implied Volatility Function (IVF) Model** Model designed so that it matches the market prices of all European options.
- Inception Profit** Profit created by selling a derivative for more than its theoretical value.
- Index Amortizing Swap** See indexed principal swap.

- Index Arbitrage** An arbitrage involving a position in the stocks comprising a stock index and a position in a futures contract on the stock index.
- Index Futures** A futures contract on a stock index or other index.
- Index Option** An option contract on a stock index or other index.
- Indexed Principal Swap** A swap where the principal declines over time. The reduction in the principal on a payment date depends on the level of interest rates.
- Initial Margin** The cash required from a futures trader at the time of the trade.
- Instantaneous Forward Rate** Forward rate for a very short period of time in the future.
- Interest Rate Cap** An option that provides a payoff when a specified interest rate is above a certain level. The interest rate is a floating rate that is reset periodically.
- Interest Rate Collar** A combination of an interest-rate cap and an interest rate floor.
- Interest Rate Derivative** A derivative whose payoffs are dependent on future interest rates.
- Interest Rate Floor** An option that provides a payoff when an interest rate is below a certain level. The interest rate is a floating rate that is reset periodically.
- Interest Rate Option** An option where the payoff is dependent on the level of interest rates.
- Interest Rate Swap** An exchange of a fixed rate of interest on a certain notional principal for a floating rate of interest on the same notional principal.
- International Swaps and Derivatives Association** Trade Association for over-the-counter derivatives and developer of master agreements used in over-the-counter contracts.
- In-the-Money Option** Either (a) a call option where the asset price is greater than the strike price or (b) a put option where the asset price is less than the strike price.
- Intrinsic Value** For a call option, this is the greater of the excess of the asset price over the strike price and zero. For a put option, it is the greater of the excess of the strike price over the asset price and zero.
- Inverted Market** A market where futures prices decrease with maturity.
- Investment Asset** An asset held by at least some individuals for investment purposes.
- IO** Interest Only. A mortgage-backed security where the holder receives only interest cash flows on the underlying mortgage pool.
- ISDA** *See* International Swaps and Derivatives Association.
- Itô Process** A stochastic process where the change in a variable during each short period of time of length Δt has a normal distribution. The mean and variance of the distribution are proportional to Δt and are not necessarily constant.
- Itô's Lemma** A result that enables the stochastic process for a function of a variable to be calculated from the stochastic process for the variable itself.
- ITraxx Europe** Portfolio of 125 investment-grade European companies.
- Jump-Diffusion Model** Model where asset price has jumps superimposed on to a diffusion process such as geometric Brownian motion.

- Jump Process** Stochastic process for a variable involving jumps in the value of the variable.
- Kurtosis** A measure of the fatness of the tails of a distribution.
- LEAPS** Long-term equity anticipation securities. These are relatively long-term options on individual stocks or stock indices.
- LIBID** London interbank bid rate. The rate bid by banks on Eurocurrency deposits (i.e., the rate at which a bank is willing to borrow from other banks).
- LIBOR** London interbank offered rate. The rate offered by banks on Eurocurrency deposits (i.e., the rate at which a bank is willing to lend to other banks).
- LIBOR Curve** LIBOR zero-coupon interest rates as a function of maturity.
- LIBOR-in-Arrears Swap** Swap where the interest paid on a date is determined by the interest rate observed on that date (not by the interest rate observed on the previous payment date).
- Limit Move** The maximum price move permitted by the exchange in a single trading session.
- Limit Order** An order that can be executed only at a specified price or one more favorable to the investor.
- Liquidity Preference Theory** A theory leading to the conclusion that forward interest rates are above expected future spot interest rates.
- Liquidity Premium** The amount that forward interest rates exceed expected future spot interest rates.
- Liquidity Risk** Risk that it will not be possible to sell a holding of a particular instrument at its theoretical price. Also, the risk that a company will not be able to borrow money to fund its assets.
- Locals** Individuals on the floor of an exchange who trade for their own account rather than for someone else.
- Lognormal Distribution** A variable has a lognormal distribution when the logarithm of the variable has a normal distribution.
- Long Hedge** A hedge involving a long futures position.
- Long Position** A position involving the purchase of an asset.
- Lookback Option** An option whose payoff is dependent on the maximum or minimum of the asset price achieved during a certain period.
- Low Discrepancy Sequence** *See* Quasi-random Sequence.
- Maintenance Margin** When the balance in a trader's margin account falls below the maintenance margin level, the trader receives a margin call requiring the account to be topped up to the initial margin level.
- Margin** The cash balance (or security deposit) required from a futures or options trader.
- Margin Call** A request for extra margin when the balance in the margin account falls below the maintenance margin level.
- Market Maker** A trader who is willing to quote both bid and offer prices for an asset.
- Market Model** A model most commonly used by traders.

- Market Price of Risk** A measure of the trade-offs investors make between risk and return.
- Market Segmentation Theory** A theory that short interest rates are determined independently of long interest rates by the market.
- Marking to Market** The practice of revaluing an instrument to reflect the current values of the relevant market variables.
- Markov Process** A stochastic process where the behavior of the variable over a short period of time depends solely on the value of the variable at the beginning of the period, not on its past history.
- Martingale** A zero drift stochastic process.
- Maturity Date** The end of the life of a contract.
- Maximum Likelihood Method** A method for choosing the values of parameters by maximizing the probability of a set of observations occurring.
- Mean Reversion** The tendency of a market variable (such as an interest rate) to revert back to some long-run average level.
- Measure** Sometimes also called a probability measure, it defines the market price of risk.
- Mezzanine Tranche** Tranche which experiences losses after equity tranche but before senior tranches.
- Modified Duration** A modification to the standard duration measure so that it more accurately describes the relationship between proportional changes in a bond price and actual changes in its yield. The modification takes account of the compounding frequency with which the yield is quoted.
- Money Market Account** An investment that is initially equal to \$1 and, at time t , increases at the very short-term risk-free interest rate prevailing at that time.
- Monte Carlo Simulation** A procedure for randomly sampling changes in market variables in order to value a derivative.
- Mortgage-Backed Security** A security that entitles the owner to a share in the cash flows realized from a pool of mortgages.
- Naked Position** A short position in a call option that is not combined with a long position in the underlying asset.
- Netting** The ability to offset contracts with positive and negative values in the event of a default by a counterparty.
- Newton–Raphson Method** An iterative procedure for solving nonlinear equations.
- NINJA** Term used to describe a person with a poor credit risk: no income, no job, no assets.
- No-Arbitrage Assumption** The assumption that there are no arbitrage opportunities in market prices.
- No-Arbitrage Interest Rate Model** A model for the behavior of interest rates that is exactly consistent with the initial term structure of interest rates.
- Nonstationary Model** A model where the volatility parameters are a function of time.
- Nonsystematic Risk** Risk that can be diversified away.

- Normal Backwardation** A situation where the futures price is below the expected future spot price.
- Normal Distribution** The standard bell-shaped distribution of statistics.
- Normal Market** A market where futures prices increase with maturity.
- Notional Principal** The principal used to calculate payments in an interest rate swap. The principal is “notional” because it is neither paid nor received.
- Numeraire** Defines the units in which security prices are measured. For example, if the price of IBM is the numeraire, all security prices are measured relative to IBM. If IBM is \$80 and a particular security price is \$50, the security price is 0.625 when IBM is the numeraire.
- Numerical Procedure** A method of valuing an option when no formula is available.
- OCC** Options Clearing Corporation. *See* Clearinghouse.
- Offer Price** *See* Ask Price.
- OIS** *See* Overnight Indexed Swap.
- Open Interest** The total number of long positions outstanding in a futures contract (equals the total number of short positions).
- Open Outcry** System of trading where traders meet on the floor of the exchange
- Option** The right to buy or sell an asset.
- Option-Adjusted Spread** The spread over the Treasury curve that makes the theoretical price of an interest rate derivative equal to the market price.
- Option Class** All options of the same type (call or put) on a particular stock.
- Option Series** All options of a certain class with the same strike price and expiration date.
- Order Book Official** *See* Board Broker.
- Out-of-the-Money Option** Either (a) a call option where the asset price is less than the strike price or (b) a put option where the asset price is greater than the strike price.
- Overnight Indexed Swap** Swap where a fixed rate for a period (e.g., 1 month) is exchanged for the geometric average of the overnight rates during the period.
- Over-the-Counter Market** A market where traders deal by phone. The traders are usually financial institutions, corporations, and fund managers.
- Package** A derivative that is a portfolio of standard calls and puts, possibly combined with a position in forward contracts and the asset itself.
- Par Value** The principal amount of a bond.
- Par Yield** The coupon on a bond that makes its price equal the principal.
- Parallel Shift** A movement in the yield curve where each point on the curve changes by the same amount.
- Parisian Option** Barrier option where the asset has to be above or below the barrier for a period of time for the option to be knocked in or out.
- Path-Dependent Option** An option whose payoff depends on the whole path followed by the underlying variable—not just its final value.

- Payoff** The cash realized by the holder of an option or other derivative at the end of its life.
- PD** Probability of default.
- Plain Vanilla** A term used to describe a standard deal.
- P-Measure** Real-world measure.
- PO** Principal Only. A mortgage-backed security where the holder receives only principal cash flows on the underlying mortgage pool.
- Poisson Process** A process describing a situation where events happen at random. The probability of an event in time Δt is $\lambda \Delta t$, where λ is the intensity of the process.
- Portfolio Immunization** Making a portfolio relatively insensitive to interest rates.
- Portfolio Insurance** Entering into trades to ensure that the value of a portfolio will not fall below a certain level.
- Position Limit** The maximum position a trader (or group of traders acting together) is allowed to hold.
- Premium** The price of an option.
- Prepayment function** A function estimating the prepayment of principal on a portfolio of mortgages in terms of other variables.
- Principal** The par or face value of a debt instrument.
- Principal Components Analysis** An analysis aimed at finding a small number of factors that describe most of the variation in a large number of correlated variables (similar to a factor analysis).
- Principal Protected Note** A product where the return earned depends on the performance of a risky asset but is guaranteed to be nonnegative, so that the investor's principal is preserved.
- Program Trading** A procedure where trades are automatically generated by a computer and transmitted to the trading floor of an exchange.
- Protective Put** A put option combined with a long position in the underlying asset.
- Pull-to-Par** The reversion of a bond's price to its par value at maturity.
- Put-Call Parity** The relationship between the price of a European call option and the price of a European put option when they have the same strike price and maturity date.
- Put Option** An option to sell an asset for a certain price by a certain date.
- Puttable Bond** A bond where the holder has the right to sell it back to the issuer at certain predetermined times for a predetermined price.
- Puttable Swap** A swap where one side has the right to terminate early.
- Q-Measure** Risk-neutral measure.
- Quanto** A derivative where the payoff is defined by variables associated with one currency but is paid in another currency.
- Quasi-random Sequences** A sequences of numbers used in a Monte Carlo simulation that are representative of alternative outcomes rather than random.
- Rainbow Option** An option whose payoff is dependent on two or more underlying variables.

- Range Forward Contract** The combination of a long call and short put or the combination of a short call and long put.
- Ratchet Cap** Interest rate cap where the cap rate applicable to an accrual period equals the rate for the previous accrual period plus a spread.
- Real Option** Option involving real (as opposed to financial) assets. Real assets include land, plant, and machinery.
- Rebalancing** The process of adjusting a trading position periodically. Usually the purpose is to maintain delta neutrality.
- Recovery Rate** Amount recovered in the event of a default as a percent of the face value.
- Reference Entity** Company for which default protection is bought in a credit default swap.
- Repo** Repurchase agreement. A procedure for borrowing money by selling securities to a counterparty and agreeing to buy them back later at a slightly higher price.
- Repo Rate** The rate of interest in a repo transaction.
- Reset Date** The date in a swap or cap or floor when the floating rate for the next period is set.
- Reversion Level** The level to which the value of a market variable (e.g., an interest rate) tends to revert.
- Rho** Rate of change of the price of a derivative with the interest rate.
- Rights Issue** An issue to existing shareholders of a security giving them the right to buy new shares at a certain price.
- Risk-Free Rate** The rate of interest that can be earned without assuming any risks.
- Risk-Neutral Valuation** The valuation of an option or other derivative assuming the world is risk neutral. Risk-neutral valuation gives the correct price for a derivative in all worlds, not just in a risk-neutral world.
- Risk-Neutral World** A world where investors are assumed to require no extra return on average for bearing risks.
- Roll Back** *See* Backwards Induction.
- Scalper** A trader who holds positions for a very short period of time.
- Scenario Analysis** An analysis of the effects of possible alternative future movements in market variables on the value of a portfolio.
- SEC** Securities and Exchange Commission.
- Securitization** Procedure for distributing the risks in a portfolio of assets.
- Settlement Price** The average of the prices that a contract trades for immediately before the bell signaling the close of trading for a day. It is used in mark-to-market calculations.
- Sharpe Ratio** Ratio of excess return over risk-free rate to standard deviation of the excess return.
- Short Hedge** A hedge where a short futures position is taken.
- Short Position** A position assumed when traders sell shares they do not own.
- Short Rate** The interest rate applying for a very short period of time.

- Short Selling** Selling in the market shares that have been borrowed from another investor.
- Short-Term Risk-Free Rate** *See* Short Rate.
- Shout Option** An option where the holder has the right to lock in a minimum value for the payoff at one time during its life.
- Simulation** *See* Monte Carlo Simulation.
- Specialist** An individual responsible for managing limit orders on some exchanges. The specialist does not make the information on outstanding limit orders available to other traders.
- Speculator** An individual who is taking a position in the market. Usually the individual is betting that the price of an asset will go up or that the price of an asset will go down.
- Spot Interest Rate** *See* Zero-Coupon Interest Rate.
- Spot Price** The price for immediate delivery.
- Spot Volatilities** The volatilities used to price a cap when a different volatility is used for each caplet.
- Spread Option** An option where the payoff is dependent on the difference between two market variables.
- Spread Transaction** A position in two or more options of the same type.
- Stack and Roll** Procedure where short-term futures contracts are rolled forward so that long-term hedges are created.
- Static Hedge** A hedge that does not have to be changed once it is initiated.
- Static Options Replication** A procedure for hedging a portfolio that involves finding another portfolio of approximately equal value on some boundary.
- Step-up Swap** A swap where the principal increases over time in a predetermined way.
- Sticky Cap** Interest rate cap where the cap rate applicable to an accrual period equals the capped rate for the previous accrual period plus a spread.
- Stochastic Process** An equation describing the probabilistic behavior of a stochastic variable.
- Stochastic Variable** A variable whose future value is uncertain.
- Stock Dividend** A dividend paid in the form of additional shares.
- Stock Index** An index monitoring the value of a portfolio of stocks.
- Stock Index Futures** Futures on a stock index.
- Stock Index Option** An option on a stock index.
- Stock Option** Option on a stock.
- Stock Split** The conversion of each existing share into more than one new share.
- Storage Costs** The costs of storing a commodity.
- Straddle** A long position in a call and a put with the same strike price.
- Strangle** A long position in a call and a put with different strike prices.
- Strap** A long position in two call options and one put option with the same strike price.

- Stressed VaR** Value at risk calculated using historical simulation from a period of stressed market conditions.
- Stress Testing** Testing of the impact of extreme market moves on the value of a portfolio.
- Strike Price** The price at which the asset may be bought or sold in an option contract (also called the exercise price).
- Strip** A long position in one call option and two put options with the same strike price.
- Strip Bonds** Zero-coupon bonds created by selling the coupons on Treasury bonds separately from the principal.
- Subprime Mortgage** Mortgage granted to borrower with a poor credit history or no credit history.
- Swap** An agreement to exchange cash flows in the future according to a prearranged formula.
- Swap Rate** The fixed rate in an interest rate swap that causes the swap to have a value of zero.
- Swaption** An option to enter into an interest rate swap where a specified fixed rate is exchanged for floating.
- Swing Option** Energy option in which the rate of consumption must be between a minimum and maximum level. There is usually a limit on the number of times the option holder can change the rate at which the energy is consumed.
- Synthetic CDO** A CDO created by selling credit default swaps.
- Synthetic Option** An option created by trading the underlying asset.
- Systematic Risk** Risk that cannot be diversified away.
- Systemic Risk** Risk that a default by one financial institution will lead to defaults by other financial institutions.
- Tailing the Hedge** A procedure for adjusting the number of futures contracts used for hedging to reflect daily settlement.
- Tail Loss** *See* Conditional Value at Risk.
- Take-and-Pay Option** *See* Swing Option.
- TED Spread** The difference between 3-month LIBOR and the 3-month T-Bill rate.
- Term Structure of Interest Rates** The relationship between interest rates and their maturities.
- Terminal Value** The value at maturity.
- Theta** The rate of change of the price of an option or other derivative with the passage of time.
- Time Decay** *See* Theta.
- Time Value** The value of an option arising from the time left to maturity (equals an option's price minus its intrinsic value).
- Timing Adjustment** Adjustment made to the forward value of a variable to allow for the timing of a payoff from a derivative.

- Total Return Swap** A swap where the return on an asset such as a bond is exchanged for LIBOR plus a spread. The return on the asset includes income such as coupons and the change in value of the asset.
- Tranche** One of several securities that have different risk attributes. Examples are the tranches of a CDO or CMO.
- Transaction Costs** The cost of carrying out a trade (commissions plus the difference between the price obtained and the midpoint of the bid–offer spread).
- Treasury Bill** A short-term non-coupon-bearing instrument issued by the government to finance its debt.
- Treasury Bond** A long-term coupon-bearing instrument issued by the government to finance its debt.
- Treasury Bond Futures** A futures contract on Treasury bonds.
- Treasury Note** *See* Treasury Bond. (Treasury notes have maturities of less than 10 years.)
- Treasury Note Futures** A futures contract on Treasury notes.
- Tree** Representation of the evolution of the value of a market variable for the purposes of valuing an option or other derivative.
- Trinomial Tree** A tree where there are three branches emanating from each node. It is used in the same way as a binomial tree for valuing derivatives.
- Triple Witching Hour** A term given to the time when stock index futures, stock index options, and options on stock index futures all expire together.
- Underlying Variable** A variable on which the price of an option or other derivative depends.
- Unsystematic Risk** *See* Nonsystematic Risk.
- Up-and-In Option** An option that comes into existence when the price of the underlying asset increases to a prespecified level.
- Up-and-Out Option** An option that ceases to exist when the price of the underlying asset increases to a prespecified level.
- Uptick** An increase in price.
- Value at Risk** A loss that will not be exceeded at some specified confidence level.
- Variance–Covariance Matrix** A matrix showing variances of, and covariances between, a number of different market variables.
- Variance-Gamma Model** A pure jump model where small jumps occur often and large jumps occur infrequently.
- Variance Rate** The square of volatility.
- Variance Reduction Procedures** Procedures for reducing the error in a Monte Carlo simulation.
- Variance Swap** Swap where the realized variance rate during a period is exchanged for a fixed variance rate. Both are applied to a notional principal.
- Variation Margin** An extra margin required to bring the balance in a margin account up to the initial margin when there is a margin call.
- Vega** The rate of change in the price of an option or other derivative with volatility.

- Vega-Neutral Portfolio** A portfolio with a vega of zero.
- Vesting Period** Period during which an option cannot be exercised.
- VIX Index** Index of the volatility of the S&P 500.
- Volatility** A measure of the uncertainty of the return realized on an asset.
- Volatility Skew** A term used to describe the volatility smile when it is nonsymmetrical.
- Volatility Smile** The variation of implied volatility with strike price.
- Volatility Surface** A table showing the variation of implied volatilities with strike price and time to maturity.
- Volatility Swap** Swap where the realized volatility during a period is exchanged for a fixed volatility. Both percentage volatilities are applied to a notional principal.
- Volatility Term Structure** The variation of implied volatility with time to maturity.
- Warrant** An option issued by a company or a financial institution. Call warrants are frequently issued by companies on their own stock.
- Waterfall** Rules determining how cash flows from the underlying portfolio are distributed to tranches.
- Weather Derivative** Derivative where the payoff depends on the weather.
- Wiener Process** A stochastic process where the change in a variable during each short period of time of length Δt has a normal distribution with a mean equal to zero and a variance equal to Δt .
- Wild Card Play** The right to deliver on a futures contract at the closing price for a period of time after the close of trading.
- Writing an Option** Selling an option.
- Yield** A return provided by an instrument.
- Yield Curve** *See* Term Structure.
- Zero-Coupon Bond** A bond that provides no coupons.
- Zero-Coupon Interest Rate** The interest rate that would be earned on a bond that provides no coupons.
- Zero-Coupon Yield Curve** A plot of the zero-coupon interest rate against time to maturity.
- Zero Curve** *See* Zero-Coupon Yield Curve.
- Zero Rate** *See* Zero-Coupon Interest Rate.



DerivaGem Software

There are a number of new features of DerivaGem. The software has been simplified by eliminating the *.dll files. Source code is included with the functions, and functions are now accessible to Mac and Linux users. CDSs and CDOs can now be valued.

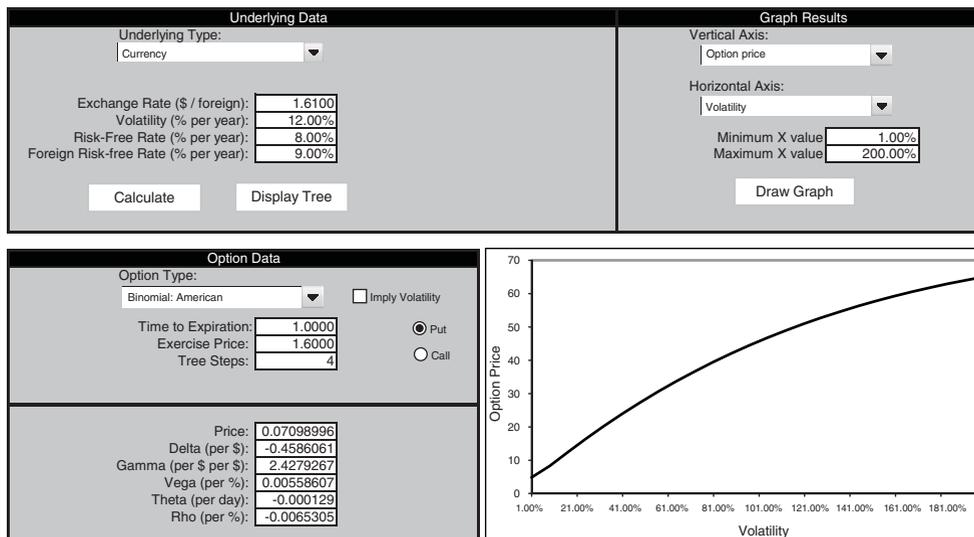
Getting Started

The most difficult part of using software is getting started. Here is a step-by-step guide to valuing an option using DerivaGem Version 2.01.

1. Put the disk that comes with this book into the CD/DVD drive on your computer. Open the Excel file DG201.xls
2. If you are using Office 2007, click on *Options* at the top of your screen (above the F column) and then click *Enable this content*. If you are not using Office 2007, make sure that the security for macros is set at medium or low. (You can do this by clicking *Tools*, followed by *Macros*, followed by *Security*.)
3. Click on the *Equity_FX_Index_Futures* worksheet tab at the bottom of the page.
4. Choose *Currency* as the Underlying Type and *Binomial American* as the Option Type. Click on the *Put* button. Leave *Implied Volatility* unchecked.
5. You are now all set to value an American put option on a currency. There are six inputs: exchange rate, volatility, domestic risk-free rate, foreign risk-free rate rate, time to expiration (years), exercise price, and time steps. Input these in cells D6, D7, D8, D9, D19, D20, and D21 as 1.61, 12%, 8%, 9%, 1.0, 1.60, and 4, respectively.
6. Hit *Enter* on your keyboard and click on *Calculate*. You will see the price of the option in cell D25 as 0.07099 and the Greek letters in cells D26 to D30. The screen you should have produced is shown on the following page.
7. Click on *Display Tree*. You will see the binomial tree used to calculate the option. This is Figure 20.6 in Chapter 20.

Next Steps

You should now have no difficulty valuing other types of option on other underlyings with this worksheet. To imply a volatility, check the *Implied Volatility* box and input the



option price in cell D25. Hit *Enter* and click on *Calculate*. The implied volatility is displayed in cell D7.

Many different charts can be displayed. To display a chart, you must first choose the variable you require on the vertical axis, the variable you require on the horizontal axis, and the range of values to be considered on the horizontal axis. Following that, you should hit *Enter* on your keyboard and click on *Draw Graph*.

Other points to note about this worksheet are:

1. For European and American equity options, up to 10 dividends on the underlying stock can be input in a table that pops up. Enter the time of each dividend (measured in years from today) in the first column and the amount of the dividend in the second column. Dividends must be entered in chronological order.
2. Up to 500 time steps can be used for the valuation of American options, but only a maximum of 10 time steps can be displayed.
3. Greek letters for all options other than standard calls and puts are calculated by perturbing the inputs, not by using analytic formulas.
4. For an Asian option the *Current Average* is the average price since inception. For a new deal (with zero time to inception), the current average is irrelevant.
5. In the case of a lookback option, *Minimum to Date* is used when a call is valued and *Maximum to Date* is used when a put is valued. For a new deal, these should be set equal to the current price of the underlying asset.
6. Interest rates are continuously compounded.

Bond Options

The general operation of the *Bond_Options* worksheet is similar to that of the *Equity_FX_Index_Futures* worksheet. The alternative models are Black's model (see Section 28.1), the normal model of the short rate (see equation (30.13)), and the lognormal model of the short rate (see equation (30.18)). The first model can be applied

only to European options. The other two can be applied to European or American options. The coupon is the rate paid per year and the frequency of payments can be selected as Quarterly, Semi-Annual or Annual. The zero-coupon yield curve is entered in the table labeled Term Structure. Enter maturities (measured in years) in the first column and the corresponding continuously compounded rates in the second column. The maturities must be in chronological order. DerivaGem assumes a piecewise linear zero curve similar to that in Figure 4.1. The strike price can be quoted (clean) or cash (dirty) (see Section 28.1). The quoted bond price, which is calculated by the software, and the strike price, which is input, are per \$100 of principal.

Caps and Swaptions

The general operation of the *Caps_and Swap_Options* worksheet is similar to that of the *Equity_FX_Index_Futures* worksheet. The worksheet is used to value interest rate caps/floors and swap options. Black's model for caps and floors is explained in Section 28.2 and Black's model for European swap options is explained in Section 28.3. The normal and lognormal short-rate models are in equations (30.13) and (30.18), respectively. The term structure of interest rates is entered in the same way as for bond options. The frequency of payments can be selected as Monthly, Quarterly, Semi-Annual, or Annual. The software calculates payment dates by working backward from the end of the life of the instrument. The initial accrual period for a cap/floor may be a nonstandard length between 0.5 and 1.5 times a normal accrual period.

CDSs

The CDS worksheet is used to calculate hazard rates from CDS spreads and vice versa. Users must input a term structure of interest rates (continuously compounded) and either a term structure of CDS spreads or a term structure of hazard rates. The initial hazard rate applies from time zero to the time specified; the second hazard rate applies from the time corresponding to the first hazard rate to the time corresponding to the second hazard rate; and so on. The hazard rates are continuously compounded, so that a hazard rate $h(t)$ at time t means that the probability of default between times t and $t + \Delta t$, conditional on no earlier default, is $h(t) \Delta t$. The calculations are carried out assuming that default can occur only at points midway between payment dates. This corresponds to the calculations for the example in Section 23.2 (the hazard rate in that example is 2% with annual compounding or 2.02% with continuous compounding).

CDOs

The CDO worksheet calculates quotes for the tranches of CDOs from tranche correlations input by the user. The attachment points and detachment points for tranches are input by the user. The quotes can be in basis points or involve an upfront payment. In the latter case, the spread in basis points is fixed and the upfront payment, as a percent of the tranche principal, is either input or implied. (For example, the fixed spread for the equity tranche of iTraxx Europe or CDX NA IG is 500 basis points.) The number of integration points (see equation (24.12)) defines the accuracy of calculations and can be left as 10 for most purposes (the maximum is 30). The software displays the expected loss as a percent of the tranche principal (ExpLoss) and the present value of expected

payments (PVPmts) at the rate of 10,000 basis points per year. The spread and upfront payment are

$$\text{ExpLoss} * 10,000/\text{PVPmts} \quad \text{and} \quad \text{ExpLoss} - (\text{Spread} * \text{PVPmts}/10,000)$$

respectively. The worksheet can be used to imply either tranche (compound) correlations or base correlations from quotes input by the user. For base correlations to be calculated, it is necessary for the first attachment point to be 0% and the detachment point for one tranche to be the attachment point for the next tranche.

How Greek Letters Are Defined

In the *Equity_FX_Index_Futures* worksheet, the Greek letters are defined as follows:

Delta: Change in option price per dollar increase in underlying asset

Gamma: Change in delta per dollar increase in underlying asset

Vega: Change in option price per 1% increase in volatility (e.g., volatility increases from 20% to 21%)

Rho: Change in option price per 1% increase in interest rate (e.g., interest increases from 5% to 6%)

Theta: Change in option price per calendar day passing.

In the *Bond_Options* and *Caps_and_Swap_Options* worksheets, the Greek letters are defined as follows:

DV01: Change in option price per 1-basis-point upward parallel shift in the zero curve

Gamma01: Change in DV01 per 1-basis-point upward parallel shift in the zero curve, multiplied by 100

Vega: Change in option price when volatility parameter increases by 1% (e.g., volatility increases from 20% to 21%).

The Applications Builder

Once you are familiar with the Options calculator (DG201.xls), you may want to start using the Application Builder. This consists of most of the functions underlying the Options Calculator with source code. It enables you to compile tables of option values, create your own charts, or develop applications. Excel users should load DG201 functions.xls and Open Office users should load Open Office DG201 functions.ods. Below are some sample applications that have been developed. They are in DG201 applications.xls and Open Office DG201 applications.ods. If any reader wishes to distribute other applications to colleagues, I would be pleased to do this (with full acknowledgements) via my website and the next release of the software.

- A. Binomial Convergence. This investigates the convergence of the binomial model in Chapters 12 and 20.
- B. Greek Letters. This provides charts showing the Greek letters in Chapter 18.
- C. Delta Hedge. This investigates the performance of delta hedging as in Tables 18.2 and 18.3.

- D.** Delta and Gamma Hedge. This investigates the performance of delta plus gamma hedging for a position in a binary option.
- E.** Value and Risk. This calculates Value at Risk for a portfolio using three different approaches.
- F.** Barrier Replication. This carries out calculations for static options replication (see Section 25.16).
- G.** Trinomial Convergence. This investigates the convergence of a trinomial tree model.

Note that E, F, and G are not included in the Open Office version of the software.

Major Exchanges Trading Futures and Options

Australian Securities Exchange (ASX)	www.asx.com.au
BM&FBOVESPA (BMF)	www.bmfbovespa.com.br
Bombay Stock Exchange (BSE)	www.bseindia.com
Boston Options Exchange (BOX)	www.bostonoptions.com
Bursa Malaysia (BM)	www.bursamalaysia.com
Chicago Board Options Exchange (CBOE)	www.cboe.com
China Financial Futures Exchange (CFFEX)	www.cffex.com.cn
CME Group	www.cmegroup.com
Dalian Commodity Exchange (DCE)	www.dce.com.cn
Eurex	www.eurexchange.com
Hong Kong Futures Exchange (HKFE)	www.hkex.com.hk
IntercontinentalExchange (ICE)	www.theice.com
International Securities Exchange (ISE)	www.iseoptions.com
Kansas City Board of Trade (KCBT)	www.kcbt.com
London Metal Exchange (LME)	www.lme.co.uk
MEFF Renta Fija and Variable, Spain	www.meff.es
Mexican Derivatives Exchange (MEXDER)	www.mexder.com
Minneapolis Grain Exchange (MGE)	www.mgex.com
Montreal Exchange (ME)	www.m-x.ca
NASDAQ OMX	www.nasdaqomx.com
National Stock Exchange, Mumbai (NSE)	www.nseindia.com
NYSE Euronext	www.nyse.com
Osaka Securities Exchange (OSE)	www.ose.or.jp
Shanghai Futures Exchange (SHFE)	www.shfe.com.cn
Singapore Exchange (SGX)	www.sgx.com
Tokyo Grain Exchange (TGE)	www.tge.or.jp
Tokyo Financial Exchange (TFX)	www.tfx.co.jp
Zhengzhou Commodity Exchange (ZCE)	www.zce.cn

There has been a great deal of consolidation of derivatives exchanges, nationally and internationally, in the last few years. The Chicago Board of Trade and the Chicago Mercantile Exchange have merged to form the CME Group, which also includes the New York Mercantile Exchange (NYMEX). Euronext and the NYSE have merged to form NYSE Euronext, which now owns the American Stock Exchange (AMEX), the Pacific Exchange (PXS), the London International Financial Futures Exchange (LIFFE), and two French exchanges. The Australian Stock Exchange and the Sydney Futures Exchange (SFE) have merged to form the Australian Securities Exchange (ASX). The IntercontinentalExchange (ICE) has acquired the New York Board of Trade (NYBOT), the International Petroleum Exchange (IPE), and the Winnipeg Commodity Exchange (WCE). Eurex, which is jointly operated by Deutsche Borse AG and SIX Swiss Exchange, has acquired the International Securities Exchange (ISE). No doubt the consolidation has been largely driven by economies of scale that lead to lower trading costs.

Table for $N(x)$ When $x \leq 0$

This table shows values of $N(x)$ for $x \leq 0$. The table should be used with interpolation. For example,

$$\begin{aligned} N(-0.1234) &= N(-0.12) - 0.34[N(-0.12) - N(-0.13)] \\ &= 0.4522 - 0.34 \times (0.4522 - 0.4483) \\ &= 0.4509 \end{aligned}$$

x	.00	.01	.02	.03	.04	.05	.06	.07	.08	.09
-0.0	0.5000	0.4960	0.4920	0.4880	0.4840	0.4801	0.4761	0.4721	0.4681	0.4641
-0.1	0.4602	0.4562	0.4522	0.4483	0.4443	0.4404	0.4364	0.4325	0.4286	0.4247
-0.2	0.4207	0.4168	0.4129	0.4090	0.4052	0.4013	0.3974	0.3936	0.3897	0.3859
-0.3	0.3821	0.3783	0.3745	0.3707	0.3669	0.3632	0.3594	0.3557	0.3520	0.3483
-0.4	0.3446	0.3409	0.3372	0.3336	0.3300	0.3264	0.3228	0.3192	0.3156	0.3121
-0.5	0.3085	0.3050	0.3015	0.2981	0.2946	0.2912	0.2877	0.2843	0.2810	0.2776
-0.6	0.2743	0.2709	0.2676	0.2643	0.2611	0.2578	0.2546	0.2514	0.2483	0.2451
-0.7	0.2420	0.2389	0.2358	0.2327	0.2296	0.2266	0.2236	0.2206	0.2177	0.2148
-0.8	0.2119	0.2090	0.2061	0.2033	0.2005	0.1977	0.1949	0.1922	0.1894	0.1867
-0.9	0.1841	0.1814	0.1788	0.1762	0.1736	0.1711	0.1685	0.1660	0.1635	0.1611
-1.0	0.1587	0.1562	0.1539	0.1515	0.1492	0.1469	0.1446	0.1423	0.1401	0.1379
-1.1	0.1357	0.1335	0.1314	0.1292	0.1271	0.1251	0.1230	0.1210	0.1190	0.1170
-1.2	0.1151	0.1131	0.1112	0.1093	0.1075	0.1056	0.1038	0.1020	0.1003	0.0985
-1.3	0.0968	0.0951	0.0934	0.0918	0.0901	0.0885	0.0869	0.0853	0.0838	0.0823
-1.4	0.0808	0.0793	0.0778	0.0764	0.0749	0.0735	0.0721	0.0708	0.0694	0.0681
-1.5	0.0668	0.0655	0.0643	0.0630	0.0618	0.0606	0.0594	0.0582	0.0571	0.0559
-1.6	0.0548	0.0537	0.0526	0.0516	0.0505	0.0495	0.0485	0.0475	0.0465	0.0455
-1.7	0.0446	0.0436	0.0427	0.0418	0.0409	0.0401	0.0392	0.0384	0.0375	0.0367
-1.8	0.0359	0.0351	0.0344	0.0336	0.0329	0.0322	0.0314	0.0307	0.0301	0.0294
-1.9	0.0287	0.0281	0.0274	0.0268	0.0262	0.0256	0.0250	0.0244	0.0239	0.0233
-2.0	0.0228	0.0222	0.0217	0.0212	0.0207	0.0202	0.0197	0.0192	0.0188	0.0183
-2.1	0.0179	0.0174	0.0170	0.0166	0.0162	0.0158	0.0154	0.0150	0.0146	0.0143
-2.2	0.0139	0.0136	0.0132	0.0129	0.0125	0.0122	0.0119	0.0116	0.0113	0.0110
-2.3	0.0107	0.0104	0.0102	0.0099	0.0096	0.0094	0.0091	0.0089	0.0087	0.0084
-2.4	0.0082	0.0080	0.0078	0.0075	0.0073	0.0071	0.0069	0.0068	0.0066	0.0064
-2.5	0.0062	0.0060	0.0059	0.0057	0.0055	0.0054	0.0052	0.0051	0.0049	0.0048
-2.6	0.0047	0.0045	0.0044	0.0043	0.0041	0.0040	0.0039	0.0038	0.0037	0.0036
-2.7	0.0035	0.0034	0.0033	0.0032	0.0031	0.0030	0.0029	0.0028	0.0027	0.0026
-2.8	0.0026	0.0025	0.0024	0.0023	0.0023	0.0022	0.0021	0.0021	0.0020	0.0019
-2.9	0.0019	0.0018	0.0018	0.0017	0.0016	0.0016	0.0015	0.0015	0.0014	0.0014
-3.0	0.0014	0.0013	0.0013	0.0012	0.0012	0.0011	0.0011	0.0011	0.0010	0.0010
-3.1	0.0010	0.0009	0.0009	0.0009	0.0008	0.0008	0.0008	0.0008	0.0007	0.0007
-3.2	0.0007	0.0007	0.0006	0.0006	0.0006	0.0006	0.0006	0.0005	0.0005	0.0005
-3.3	0.0005	0.0005	0.0005	0.0004	0.0004	0.0004	0.0004	0.0004	0.0004	0.0003
-3.4	0.0003	0.0003	0.0003	0.0003	0.0003	0.0003	0.0003	0.0003	0.0003	0.0002
-3.5	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002
-3.6	0.0002	0.0002	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001
-3.7	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001
-3.8	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001
-3.9	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
-4.0	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table for $N(x)$ When $x \geq 0$

This table shows values of $N(x)$ for $x \geq 0$. The table should be used with interpolation. For example,

$$\begin{aligned} N(0.6278) &= N(0.62) + 0.78[N(0.63) - N(0.62)] \\ &= 0.7324 + 0.78 \times (0.7357 - 0.7324) \\ &= 0.7350 \end{aligned}$$

x	.00	.01	.02	.03	.04	.05	.06	.07	.08	.09
0.0	0.5000	0.5040	0.5080	0.5120	0.5160	0.5199	0.5239	0.5279	0.5319	0.5359
0.1	0.5398	0.5438	0.5478	0.5517	0.5557	0.5596	0.5636	0.5675	0.5714	0.5753
0.2	0.5793	0.5832	0.5871	0.5910	0.5948	0.5987	0.6026	0.6064	0.6103	0.6141
0.3	0.6179	0.6217	0.6255	0.6293	0.6331	0.6368	0.6406	0.6443	0.6480	0.6517
0.4	0.6554	0.6591	0.6628	0.6664	0.6700	0.6736	0.6772	0.6808	0.6844	0.6879
0.5	0.6915	0.6950	0.6985	0.7019	0.7054	0.7088	0.7123	0.7157	0.7190	0.7224
0.6	0.7257	0.7291	0.7324	0.7357	0.7389	0.7422	0.7454	0.7486	0.7517	0.7549
0.7	0.7580	0.7611	0.7642	0.7673	0.7704	0.7734	0.7764	0.7794	0.7823	0.7852
0.8	0.7881	0.7910	0.7939	0.7967	0.7995	0.8023	0.8051	0.8078	0.8106	0.8133
0.9	0.8159	0.8186	0.8212	0.8238	0.8264	0.8289	0.8315	0.8340	0.8365	0.8389
1.0	0.8413	0.8438	0.8461	0.8485	0.8508	0.8531	0.8554	0.8577	0.8599	0.8621
1.1	0.8643	0.8665	0.8686	0.8708	0.8729	0.8749	0.8770	0.8790	0.8810	0.8830
1.2	0.8849	0.8869	0.8888	0.8907	0.8925	0.8944	0.8962	0.8980	0.8997	0.9015
1.3	0.9032	0.9049	0.9066	0.9082	0.9099	0.9115	0.9131	0.9147	0.9162	0.9177
1.4	0.9192	0.9207	0.9222	0.9236	0.9251	0.9265	0.9279	0.9292	0.9306	0.9319
1.5	0.9332	0.9345	0.9357	0.9370	0.9382	0.9394	0.9406	0.9418	0.9429	0.9441
1.6	0.9452	0.9463	0.9474	0.9484	0.9495	0.9505	0.9515	0.9525	0.9535	0.9545
1.7	0.9554	0.9564	0.9573	0.9582	0.9591	0.9599	0.9608	0.9616	0.9625	0.9633
1.8	0.9641	0.9649	0.9656	0.9664	0.9671	0.9678	0.9686	0.9693	0.9699	0.9706
1.9	0.9713	0.9719	0.9726	0.9732	0.9738	0.9744	0.9750	0.9756	0.9761	0.9767
2.0	0.9772	0.9778	0.9783	0.9788	0.9793	0.9798	0.9803	0.9808	0.9812	0.9817
2.1	0.9821	0.9826	0.9830	0.9834	0.9838	0.9842	0.9846	0.9850	0.9854	0.9857
2.2	0.9861	0.9864	0.9868	0.9871	0.9875	0.9878	0.9881	0.9884	0.9887	0.9890
2.3	0.9893	0.9896	0.9898	0.9901	0.9904	0.9906	0.9909	0.9911	0.9913	0.9916
2.4	0.9918	0.9920	0.9922	0.9925	0.9927	0.9929	0.9931	0.9932	0.9934	0.9936
2.5	0.9938	0.9940	0.9941	0.9943	0.9945	0.9946	0.9948	0.9949	0.9951	0.9952
2.6	0.9953	0.9955	0.9956	0.9957	0.9959	0.9960	0.9961	0.9962	0.9963	0.9964
2.7	0.9965	0.9966	0.9967	0.9968	0.9969	0.9970	0.9971	0.9972	0.9973	0.9974
2.8	0.9974	0.9975	0.9976	0.9977	0.9977	0.9978	0.9979	0.9979	0.9980	0.9981
2.9	0.9981	0.9982	0.9982	0.9983	0.9984	0.9984	0.9985	0.9985	0.9986	0.9986
3.0	0.9986	0.9987	0.9987	0.9988	0.9988	0.9989	0.9989	0.9989	0.9990	0.9990
3.1	0.9990	0.9991	0.9991	0.9991	0.9992	0.9992	0.9992	0.9992	0.9993	0.9993
3.2	0.9993	0.9993	0.9994	0.9994	0.9994	0.9994	0.9994	0.9995	0.9995	0.9995
3.3	0.9995	0.9995	0.9995	0.9996	0.9996	0.9996	0.9996	0.9996	0.9996	0.9997
3.4	0.9997	0.9997	0.9997	0.9997	0.9997	0.9997	0.9997	0.9997	0.9997	0.9998
3.5	0.9998	0.9998	0.9998	0.9998	0.9998	0.9998	0.9998	0.9998	0.9998	0.9998
3.6	0.9998	0.9998	0.9999	0.9999	0.9999	0.9999	0.9999	0.9999	0.9999	0.9999
3.7	0.9999	0.9999	0.9999	0.9999	0.9999	0.9999	0.9999	0.9999	0.9999	0.9999
3.8	0.9999	0.9999	0.9999	0.9999	0.9999	0.9999	0.9999	0.9999	0.9999	0.9999
3.9	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
4.0	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000

- Cox, D. R., 294
 Cox, J. C., 125, 273, 325, 429, 467, 600, 626, 645, 685, 705, 710
 Culp, C., 68
 Cumby, R., 517
 Daghli, T., 417, 421
 Das, S., 569, 571
 Dattatreya, R. E., 176
 Delbaen, F., 472, 494
 Demeterfi, K., 589, 594
 Derman, E., 415, 421, 589, 591, 594, 607, 626, 693, 710
 Detemple, J., 231
 Dick-Nielsen, J., 529
 Dowd, K., 494
 Duan, J.-C., 606, 626
 Dufey, G., 358
 Duffie, D., 32, 42, 145, 495, 543, 645, 731
 Dunbar, N., 790
 Dupire, B., 607, 626
 Eber, J.-M., 472, 494
 Ederington, L. H., 68, 250, 421
 Edwards, D. W., 762
 Edwards, F. R., 68
 Embrechts, P., 495
 Engle, R. F., 500, 502, 506, 514, 517
 Ergener, D., 591, 594
 Eydeland, A., 757, 762
 Fabozzi, F. J., 97, 145
 Fama, E. F., 307, 324, 325
 Feldhütter, P., 529
 Feller, W., 294
 Figlewski, S., 445, 467, 517, 617, 626
 Finger, C. C., 543
 Flannery, B. P., 455, 466, 505, 708
 Flavell, R., 176
 Freed, L., 569
 French, K. R., 307, 325
 Froot, K. A., 763
 Frye, J., 491, 495
 Gao, B., 445, 467, 617, 626
 Garman, M. B., 358, 582, 645
 Gastineau, G. L., 43
 Gatarek, D., 718, 724, 731
 Gatto, M. A., 582, 594
 Geczy, C., 68
 Geman, H., 757, 762, 763
 Geng, G., 569
 Geske, R., 322, 577, 594
 Ghon, R. S., 125
 Gibson, R., 757, 762
 Giddy, I. H., 358
 Glasserman, P., 467, 581, 584
 Goldman, B., 582, 594
 Gonedes, N., 324
 Gorton, G., 192
 Grabbe, J. O., 358
 Graham, J. R., 69
 Gregory, J., 543, 571
 Grinblatt, M., 97, 145
 Guan, W., 421
 Guay, W. R., 343
 Gupta, A., 176
 Hanley, M., 763
 Hanly, J., 68
 Harrison, J. M., 645
 Hasbrook, J., 517
 Haushalter, G. D., 69
 Hayt, G. S., 68
 Heath, D., 472, 494, 715, 731
 Hendricks, D., 495
 Heron, R., 341, 343
 Heston, S. L., 606, 626
 Hicks, J. R., 121
 Hilliard, J. E., 374
 Ho, T. S. Y., 140, 690, 710
 Hopper, G., 495
 Hoskins, W., 145
 Hotta, K., 176
 Hua, P., 495
 Huddart, S., 343
 Hull, J. C., 43, 189, 190, 192, 335, 339, 343, 358, 394, 395, 412, 417, 421, 440, 453, 464, 467, 472, 490, 495, 526, 528, 541, 543, 557, 569, 570, 571, 606, 608, 611, 619, 620, 626, 691, 694, 698, 703, 705, 710, 711, 723, 724, 727, 731
 Hunter, R., 562
 Iben, B., 677
 Ingersoll, J. E., 125, 608, 645, 685, 703, 710
 Inui, K., 731
 Itô, K., 291
 Jackson, P., 495
 Jackwerth, J. C., 414, 421
 Jain, G., 467
 Jamshidian, F., 489, 495, 677, 718, 731
 Jarrow, R. A., 125, 357, 715, 731
 Jeanblanc-Picque, M., 581
 Jegadeesh, N., 145
 Jermakyan, M., 626
 Jones, F. J., 43
 Jorion, P., 43, 97, 358, 495, 790
 Ju, X., 790
 Kamal, E., 589, 594
 Kan, R., 731
 Kane, A., 517
 Kane, E. J., 125
 Kani, I., 591, 594, 607, 626
 Kapadia, N., 569
 Karasinski, P., 693, 710

- Karlin, S., 294
 Kawaller, I. G., 43
 Kealhofer, S., 543
 Kentwell, G., 581
 Keynes, J. M., 121
 Keys, B. J., 186, 192
 Kijima, M., 711, 731
 Kleinstein, A. D., 97
 Kluppelberg, C., 495
 Koch, P. D., 43
 Kohlhagen, S. W., 358
 Koller, T., 777
 Kon, S. J., 324
 Kou, S. G., 581, 584
 Kreps, D. M., 645
 Krinsman, A. N., 192
 Krzysztof, W., 762
 Kulatilaka, N., 776
 Kwok, Y. K., 583
 Lando, D., 529
 Lang, M., 343
 Lau, S. H., 615, 626
 Laurent, J.-P., 571
 Lee, R., 589, 594
 Lee, S.-B., 140, 690, 710
 Li, A., 711
 Li, D. X., 543, 571
 Lie, E., 341, 343
 Lieu, D., 373
 Litzenberger, R. H., 176, 424, 761, 763
 Ljung, G. M., 509
 Longin, F. M., 495
 Longstaff, F. A., 97, 621, 626, 686, 710, 723,
 728, 776
 Lowenstein, R., 43
 McDonald, R., 762
 McMillan, L. G., 211, 250
 McNeil, A. J., 495
 Madan, D. B., 603, 605, 626
 Margrabe, W., 586, 594
 Markowitz, H., 481
 Marshall, C., 495
 Marston, R. C., 68
 Maude, D. J., 495
 Melick, W. R., 421
 Mello, A. S., 69, 176
 Merton, R. C., 18, 224, 231, 299, 325, 351,
 358, 530, 543, 601, 626
 Mezrich, J., 506, 514, 517
 Mian, A., 192
 Mikosch, T., 495
 Milevsky, M. A., 594
 Miller, H. D., 294
 Miller, M. H., 18, 68
 Miltersen, K. R., 374, 718, 731
 Minton, B. A., 68, 176
 Mintz, D., 451
 Moon, M., 770, 777
 Morton, A., 715, 731
 Mukherjee, T., 186, 192
 Murrin, J., 777
 Musiela, M., 718, 724, 731
 Nagayama, I., 711
 Natenberg, S., 417
 Neftci, S., 294, 495
 Nelson, D., 502, 517
 Ness, A., 58
 Neuberger, A. J., 69
 Ng, V., 502, 517
 Noh, J., 517
 Oldfield, G. S., 125
 Pan, J., 495
 Parsons, J. E., 69
 Pascutti, M., 176
 Pearson, N., 790
 Perraudin, W., 495
 Persaud, A. D., 790
 Petersen, M. A., 69
 Pilipovic, D., 585
 Pindyck, R. S., 125
 Piterbarg, V., 731
 Pliska, S. R., 645
 Posner, S. E., 594
 Predescu, M., 526, 528, 543, 569
 Press, W. H., 455, 466, 505, 708
 Rebonato, R., 626, 711, 731
 Reiner, E., 594, 677
 Reis, J., 374
 Reisch, D., 421
 Remolona, E. M., 529
 Renault, O., 569
 Rendleman, R. J., 69, 142, 250, 273, 467, 683
 Rennie, A., 645
 Resti, A., 524, 543
 Reynolds, C. E., 761, 763
 Rich, D., 495, 746
 Richard, S., 125
 Richardson, M., 325, 495
 Ritchken, P., 594, 627, 711
 Robinson, F. L., 97
 Roll, R., 307, 322, 325
 Ronn, A. G., 250
 Ronn, E. I., 250
 Ross, S. A., 125, 273, 325, 429, 467, 600, 626,
 635, 645, 685, 703, 710
 Routledge, B. R., 125
 Rubinstein, M., 273, 343, 414, 416, 421, 429,
 467, 577, 587, 594, 607, 620, 627
 Sandmann, K., 718, 731
 Sandor, R. L., 763

- Sankarasubramanian, L., 594, 711
Santa-Clara, P., 728
Scholes, M., 224, 231, 299, 325
Schönbucher, P. J., 571
Schrand, C., 68
Schwartz, E. S., 374, 621, 626, 686, 710, 728,
757, 762, 770, 776, 777
Scott, L., 373
Seppi, D. J., 125
Serfaty-de Medeiros, K., 68
Seru, A., 186, 192
Sevigny, A., 569
Shapiro, A., 494
Shreve, S. E., 273, 294
Shumway, T., 273
Sidenius, J., 569, 571
Siegel, M., 495
Singleton, K., 543
Sironi, A., 524, 543
Smith, C. W., 69, 325
Smith, D. J., 43, 176, 745, 746
Smith, T., 325
Sobol', I. M., 455
Sondermann, D., 718, 731
Sorkin, A. R., 790
Sosin, H., 582, 594
Spatt, C. S., 125
Stigum, M., 97
Stoll, H. R., 231
Strickland, C., 466, 594, 762
Stulz, R. M., 69, 594
Stutzer, M., 627
Subrahmanyam, M. G., 176
Sufi, A., 192
Sun, T., 176
Sundaresan, S., 125, 176
Suo, W., 417, 421, 606, 608, 626
Taleb, N. N., 404
Tavakoli, J. M., 571
Taylor, H. M., 294
Taylor, S. J., 421
Tett, G., 790
Teukolsky, S. A., 455, 466, 505, 708
Teweles, R. J., 43
Thiagarajan, S. R., 69
Thomas, C. P., 421
Thomson, R., 790
Tilley, J. A., 621, 627
Titman, S., 176
Todd, R., 43
Tompkins, R., 417
Toy, W., 693, 710
Trevor, R., 627
Trigeorgis, L., 777
Tufano, P., 69
Turnbull, S. M., 585, 594, 731
Vasicek, O. A., 540, 543, 684, 710
Vetterling, W. T., 455, 466, 505, 708
Viciera, L. M., 68
Vig, V., 186, 192
Vijh, A. M., 594
Viswanathan, R., 626
Wakeman, L. M., 585, 594
Wang, C., 176
Wei, J., 762
Weston, J., 68
Whaley, R., 322
White, A., 189, 192, 335, 339, 343, 394, 395,
412, 440, 453, 464, 467, 490, 495, 526, 528,
543, 557, 569, 570, 571, 606, 611, 619, 620,
626, 691, 694, 698, 703, 705, 710, 711, 723,
724, 727, 731
Whitelaw, R., 495
Wiggins, J. B., 250
Wilmott, P., 467, 495, 688, 710
Wong, H. Y., 583
Wystup, U., 421
Xu, X., 421
Yermack, D., 341, 343
Yor, M., 581
Zhang, P. G., 790
Zhu, H., 32, 42
Zhu, Y., 489, 495
Zimmerman, T., 192
Zingales, L., 18
Zou, J., 589, 594



Subject Index

References to items in the Glossary of Terms are **bolded**

- Abandonment option, 769, 773
- ABS, *see* Asset-backed security.
- ABS CDO, 183–184, **791**
- Accounting, 39, 335–336
- Accrual fraction, 658, 662
- Accrual swap, 174, 742–743, **791**
- Accrued interest, 131, 133, **791**
- Adaptive Mesh Model, 445, **791**
 - valuing Barrier option, 617–618
- Add-up basket credit default swap, 558
- Adjustable rate mortgages, 185
- Agency costs, 189–190, 334, **791**
- Agency mortgage-backed securities, *see* Mortgage-backed security.
- Agricultural products, 748–749
- AIG, 32, 550
- Algorithmic trading, 2
- Allied Irish Bank, 779
- Allied Lyons, 781
- Amaranth, 779, 780
- Amazon.com, valuation, 770
- American option, 7, 194, **791**
 - analytic approximation to prices, 427
 - binomial tree, 253, 427–446
 - Black's approximation, 322–323
 - dividend, effect of, 229–230, 321–323
 - early exercise, 196, 225–227, 321–323
 - future options compared to spot options, 372
 - Monte Carlo simulation and, 621–625
 - nonstandard, 575
 - option on a dividend-paying stock, 229–230, 321–323
 - option on a non-dividend-paying stock, 225–226
 - put–call relationship, 224, 230
- Amortizing swap, 174, 733, **791**
- Analytic result, **791**
- Antithetic variable technique, variance reduction procedure, 453
- Arbitrage, **791**
- Arbitrage pricing theory (APT), 635
- Arbitrageurs, 15, 15–16, **791**
- ARCH model, 500
- Asian option, 584–586, 611, **791**
- Ask price, **791**
- Asked price, **791**
- Asset-backed security, 559, **791**
- Asset–liability management (ALM), 95, 144
- Asset-or-nothing call option, 582, **791**
- Asset-or-nothing put option, 582, **791**
- Asset swap, 526, **791**
- Assigned investor, 207
- Asymmetric information, 556
- As-you-like-it option, 578, **792**
- At-the-money option, 201, **792**
 - alternative definitions, 416
 - delta, 383
 - gamma, 392
 - theta, 388
- Attachment point, 563
- Average price call, 584
- Average price option, **792**
- Average price put option, 584
- Average strike call, 586
- Average strike option, **792**
- Average strike put, 586
- Backdating, 341–342, **792**
- Back office, 784
- Back testing, 490, **792**
- Backward difference approximation, 456
- Backwards induction, **792**
- Bankers Trust (BT), 744–745, 785
- Barings Bank, 17, 779, 782, 784

- Barrier options, 579–581, 591–593, 615–618, **792**
 inner barrier, 615–616
 outer barrier, 615–616
 valuing, using adaptive mesh model, 617–618
- Base correlation, 567, **792**
- Basel I, II, III, 190–191
- Basel Committee, 472, **792**
- Basis, 52–56, **792**
- Basis point, 90, **792**
- Basis risk, **792**
 hedging and, 52–56
- Basis swap, 734, **792**
- Basket credit default swap, 558, **792**
 add-up basket credit default swap, 558
 first-to-default basket credit default swap, 558
 kth-to-default basket credit default swap, 558
 role of correlation, 561–562
 valuation, 566
- Basket option, 587, **792**
- Bear spread, 240–241, **792**
- Bearish calendar spread, 245
- Bermudan option, 575, **792**
- Bermudan swap option, 727–728
- Beta, 62–65, 346–347, **792**
 changing, 64–65
- BGM model, *see* Brace–Gatarek–Musielka model.
- Bid, 3, 203
- Bid–ask spread, **792**
- Bid–offer spread, 153, 203, 203, **792**
- Bid price, **792**
- Binary credit default swap, 554, **792**
- Binary option, 581–582, **792**
- Binomial correlation measure, 537
- Binomial model, 253–273, 427–428, **793**
- Binomial tree, 253–273, 427–446, **793**
 alternative procedures for constructing, 442–445
 American options, 263–264, 430–433
 control variate technique, 440–442
 defined, 253–254
 delta and, 264–265
 deriving the Black–Scholes–Merton formula from, 276–279
 dividend-paying stocks, 437–440
 employee stock options, 338
 European options examples, 253–257
 futures option, 367–369
 matching volatility, 265–267
 no-arbitrage argument, 253–257
 non-dividend-paying stock, 427–434
 one-step, 253–257
 options on index, currency, and futures, 269–272, 435–437
 risk-neutral valuation and, 257–259
 stock paying a known dividend yield, 269–270, 437
 time-dependent interest rates and, 445–446
 two-step, 259–264
- Bivariate normal distribution, 578, **793**
- Black Monday, 114. *See also* October 19, 1987.
- Black–Karasinski model, 693–694
 using Hull–White tree-building procedure for, 704–705
- Black’s approximation, **793**
 American call option, 322–323
- Black–Derman–Toy model, 693
- Black’s model, 370–372, 641, 663, **793**
 bond options and, 649–652
 caps and floors and, 654–658
 forward risk-neutral valuation and, 641
 generalization of, 663
 pricing European options on spot, 371–372, 641
 swaptions and, 659–661
- Black–Scholes–Merton and volatility smile, equity option, 414–415
 foreign currency options, 411–413
 jump in asset price anticipated, 419–420
- Black–Scholes–Merton model, 299–324, **793**
 and Monte Carlo simulation, 448–449
 cumulative normal distribution function, 315
 delta and, 381–382
 deriving from binomial trees, 276–279
 deriving using risk-neutral valuation, 329–331
 differential equation, 307–311
 dividend, 320–323
 European option on non-dividend-paying stock, 313–316
 expected return, 302–303
 implied volatility, 318–320. *See also* Volatility smile.
 known dividend yield, 350–352
 pricing formulas, 313
 risk-neutral valuation and, 311–313, 314
 volatility, 303–307, 409–421
- BM&F Bovespa, 21
- Board order, 37
- Bond option, 648–652, **793**
 embedded, 648–649
 European, 649–652
 forward risk-neutral valuation and Black’s model, 641
 on coupon-bearing bonds, 695

- on zero-coupon bond, 694–695
 - tree for American bond options, 707–708
 - yield volatilities, 652
- Bond price process,
 - Cox, Ingersoll, Ross, 687
 - BGM, 719
 - HJM, 716
 - Vasicek, 687
- Bond, pricing, 80–82
- Bond valuation
 - Cox, Ross, Rubinstein model, 685
 - general short-rate model, 683
 - Vasicek model, 684–685
- Bond yield, 81, **793**
- Bond yield volatilities, 652
- Bootstrap method, 82–84, **793**
- Boston option, 575, **793**
- Boston Options Exchange, 199
- Bottom straddle, 247
- Bottom vertical combination, 247
- Bounds for options, 218–221, 226, 227–228, 350–351, 366–367
- Boundary conditions, 310
- Box spread, 241, 244, **793**
- Brace–Gatarek–Musielka model, 718–728
- Break forward, 575, **793**
- British Local Authority, 781
- Brownian motion, 282, **793**
- Bull spread, 238–239, **793**
- Bullish calendar spread, 245
- Business day conventions, 155
- Business valuation, 769
- Butterfly spread, 242–243, 249, 424–425, **793**
- Buying on margin, 205
- Calendar days vs. trading days, 306–307
- Calendar spread, 244–245, **793**
- Calibrating instruments, 707
- Calibration, 707–709, 726–727, **793**
- Call option, 7, 194–195, **793**
- Callable bond, 648–649, **793**
- Cancelable compounding swaps, 743
- Cancelable forward, 575
- Cancelable swap, 743, **793**
- Cap, interest rate, 653–656, **793**. *See also*
 - Interest rate caps and floors.
 - flexi cap, 724, 799
 - ratchet cap, 723, **806**
 - sticky cap, 723, 807
- Capital asset pricing model, 62–65, 122, 635, **794**
 - relation to capital investment appraisal, 766
 - relation to market price of risk, 632, 768–769
- Cap rate, 653, **794**
- Capital investment appraisal, 765–766
 - NPV vs. real options approach, 766
- Caplet, interest rate, 654, **794**. *See also* Cap.
- CAPM, 73–74
- Case–Shiller index, 184, **794**
- Cash CDO, 560
- Cash-flow mapping, 483–484, 489–490, **794**
- Cash-or-nothing call option, 581, **794**
- Cash-or-nothing put option, 582, **794**
- Cash price, bond and Treasury bill, 131, 650, 707. *See also* Dirty price.
- Cash settlement, 36–37, **794**
- CAT bond, catastrophic bond, 759, **794**
- CDD, *see* Cooling degree days.
- CDO, *see* Collateralized debt obligation.
- CDO squared, **794**
- CDS, *see* Credit default swap.
- CDS–bond basis, 551
- CDX NA IG, 555, 561, **794**
- Central clearing, 32
- Central clearing party, 32, **794**
- Changing the measure/numeraire, 267, 643–644
- Cheapest-to-deliver bond, **794**
 - credit default swap, 551
 - futures contract, 134–135
- Chicago Board of Trade (CBOT), 2, 7
- Chicago Board Options Exchange (CBOE), 2, 8, 199, 211, 345
- Chicago Mercantile Exchange (CME), 2, 7, 758
- Cholesky decomposition, 450, **794**
- Chooser option, 578–579, **794**
- Citigroup, 788
- Claim in event of default, 523
- Class of options. *See* Option class.
- Clean price, bond, 650, **794**. *See also* Quoted price.
- Clearing house, **794**
 - futures, 29–30
 - options, 206, **805**
 - OTC markets, 32–33
 - swaps, 172, 550
- Clearing margin, 30, **794**
- Cliquet options, 577, **794**
- CME Group, 2, 7, 22
- CMO, *see* Collateralized mortgage obligation.
- Collar, interest rate, 654, **794**
- Collateralization, 31, 535–536, **794**
- Collateralized debt obligation (CDO), 559–570, **795**
 - cash, 560
 - synthetic, 560
- Collateralized mortgage obligation (CMO), 729, **795**

- Combination, option trading strategy, 246–248, **795**
- Commission, stock option, 204–205
- Commodity,
 - agricultural, 748–749
 - crude oil, 750
 - electricity, 751–752
 - metals, 749–750
 - natural gas, 750–751
 - recycling, 749
- Commodity Futures Trading Commission (CFTC), 38, **795**
- Commodity price, 769
 - mean reversion, 749, 753–755
 - modeling, 752–757
 - seasonality, 749, 755–757
 - trinomial tree, 752–757
- Commodity swap, 175, 744, **795**
- Comparative-advantage argument,
 - currency swap, 166–168
 - interest rate swap, 156–158
- Compound correlation, 566–567, **795**
- Compound option, 577–578, **795**
- Compounding frequency, 77–80, **795**
- Compounding swap, 174, 735–736, **795**
- Conditional default probabilities, 522
- Conditional value at risk (C-VaR), 472, **795**
- Confidence level (VaR), 471
- Confirmation, 155, 734, 735, 741, **795**
- Constant elasticity of variance (CEV) model, 600–601, **795**
- Constant maturity swap (CMS), 174, 738–739, **795**
- Constant maturity Treasury swap (CMT), 174, 739, **795**
- Constructive sale, 208–209
- Consumption asset, 101, **795**
- Contagion, *see* Credit contagion.
- Contango, 123, **795**
- Continuous compounding, 79–80, **795**
- Continuous-time stochastic process, 280, 281–286
- Continuous variable, 280
- Contraction option, 771
- Contract size, 25
- Control areas, electricity-producing region, 751
- Control variate technique, variance reduction procedure, 440–442, 440–442, **795**
- Convenience yield, 119–120, **795**
- Convergence arbitrage, 31, 785
- Conversion factor, 133–134, **795**
- Conversion ratio, 608
- Convertible bond, 209, 608–611, **795**
- Convexity, bonds, 92–93, **795**
 - Convexity adjustment, 668–672, 681, **796**
 - Eurodollar futures, 140–141
 - interest rates, 668–672
 - swap rates, 671–672
 - Cooling degree days (CDD), 758, **794**
 - Copula, 538–540, 540, **796**
 - Corner the market, 39
 - Cornish–Fisher expansion, 488, **796**
 - Correlated stochastic processes, 290–291
 - Correlation,
 - default, 536–540, 561–562
 - monitoring, 512–515
 - Correlation matrix, 482–483, 515–516
 - Correlation smiles, 567
 - Cost of carry, 120, **796**
 - Counterparty, **796**
 - default risk, 532–534
 - Coupon, **796**
 - Covariance, 512–515, **796**
 - Covariance matrix, 482–483, 515–516, **796**
 - consistency condition, 514
 - Covered call, 206, 236, **796**
 - Covered position, 378
 - Crank–Nicolson procedure, 465
 - Cox–Ingersoll–Ross (CIR) interest rate model, 685–689
 - Cox–Ross–Rubinstein model, 253, 266, 427–442
 - compared with Vasicek, 687
 - Crashophobia, 416, **796**
 - Credit contagion, 505, **796**
 - Credit crisis, 180–192
 - Credit default swap (CDS), 548–555, **796**
 - basket CDS, 558
 - bond yields and, 550
 - cheapest-to-deliver bond, 551
 - fixed coupons, 556
 - forwards and options on, 557–558
 - recovery rate and, 554–555
 - spread, 549
 - valuation of, 551–554
 - Credit derivatives, 547–571, **796**
 - alternative models, 569–570
 - Credit event, 548
 - Credit index, 555–556, 561, **796**
 - Credit rating, 521, **796**
 - Credit ratings transition matrix, 542, **796**
 - Credit risk, 4, 30, 521–543, **796**
 - comparison of default probability estimates, 526–530
 - credit ratings and, 521
 - credit value at risk (credit VaR), 540–542
 - default correlation and, 536–540
 - derivatives transactions and, 531–536

- estimating default probabilities from bond prices and, 524–526
- estimating default probabilities from equity prices and, 530–531
- historical default probabilities and, 522–523
- interest rate and, 75
- mitigation, 534–536
- recovery rates and, 523–524
- swaps and, 171–172
- Credit spread option, 557–558, **796**
- Credit value adjustment (CVA), 532, **796**
- Credit value at risk, 540–542, **796**
- CreditMetrics, 541–542, **796**
- Cross-currency derivative, 674
- Cross-currency swap, 174, 737
- Cross gamma, 488
- Cross hedging, 56–60, **796**
- Crude oil derivatives, 750
- Cumulative distribution function, **796**
- Cumulative normal distribution function, 315
- Currency forward and futures, 114–117
- Currency option, 199, 347–350, 355–356
 - early exercise, 357
 - implied distribution and lognormal distribution, 412–413
 - valuation, binomial tree, 271, 435–437
 - valuation, Black–Scholes–Merton, 355–356
 - volatility smile, 411–413
- Currency swap, 165–171, 736–737, **797**
 - comparative-advantage argument, 166–168
 - to transform liabilities and assets, 166
 - valuation of, 168–171
- Curvature, 389
- CVA, *see* Credit value adjustment.
- Cylinder option, 575
- Daiwa Bank, 779
- Daily settlement, 27
- Day count conventions, 129–130, 154–155, **797**
- Day order, 38
- Day trade, 29, **797**
- Day trader, 37
- Default correlation, 536–540, **797**
 - binomial correlation measure, 537
 - factors to define correlation structure, 539–540
 - Gaussian copula model for time to default, 538–540
 - implied, 566–568
 - reduced form models, 537–538
 - structural models, 537–538
- Default intensity, 522, **797**
- Default probabilities, historical, 522–523
- Default probability, *see* Probability of default.
- Default probability density, **797**
- Default risk, 521–543
 - counterparty, adjusting derivatives' valuation, 532–534
- Deferred payment option, 575, **797**
- Deferred swap, 659, **797**
- Delivery, 24, 25, 36–37, 121
- Delivery price, **797**
- Delta, 264–265, 380–387, 398, 417–418, **797**
 - estimating, using binomial tree, 433–434
 - European options, 382–383
 - for a portfolio, 386
 - forward contract, 398–399
 - futures contract, 399
 - interest rate derivatives, 663–664
 - relationship with theta and gamma, 392–393
- Delta hedging, 264–265, 380–387, **797**
 - dynamic aspects, 383–386
 - exotic options, 591
 - impact of stochastic volatility on delta hedging, 607
 - performance measure, 380, 385
 - transaction cost, 387
- Delta-neutral portfolio, 386, 390, **797**
- DerivaGem, 269, 658, 707, **797**, 812–816
- Derivative, defined, 1, **797**
- Detachment point, 563
- Deterministic variable, **797**
- Deutsche Bank, 785
- Diagonal spread, 245–246, **797**
- Differential equation for derivative,
 - constant dividend, 352
 - no dividends, 307–311
 - on futures, 370
- Differential swap (diff swap), 174, 740, **797**
- Diffusion model, 600
- Diffusion process, **797**
- Dilution, 317, 340
- Dirty price, bond, 650, **797**. *See also* Cash price.
- Discount bond, **797**
- Discount broker, 204
- Discount instrument, **797**
- Discount rate, Treasury bill, 131
- Discrete-time stochastic process, 280, 287–289
- Discrete variable, 280
- Discretionary order, 38
- Diversification, 480–481, 783, **797**
- Dividend, 320–323, 437–440, **797**
 - American call option valuation, using Black–Scholes–Merton model, 321–323
 - binomial model for stocks paying dividends, 437–440
 - bounds of option prices, 229
 - effect on option price, 217–218

- European option valuation, using Black–Scholes–Merton model, 320–321
- stock option and, 201–202, 217–218, 320–323
- stock prices and, 201–202, 217
- stock splits and, 201–202
- Dividend yield, **797**
 - binomial tree and, 269–270, 437
 - implied, 354–355
- Dollar duration, 92, **797**
- Double t copula, 569
- Dow Jones Industrial Average (DJX), 61, 199
 - options, 345
- Down-and-in call, 579
- Down-and-in option, **797**
- Down-and-in put, 580
- Down-and-out call, 579
- Down-and-out option, **798**
- Down-and-out put, 580
- Downgrade trigger, 536, 537, **798**
- Drift rate, 284, **798**
- Duration, 89–92, **798**
 - bond, 89
 - bond portfolio, 92
 - modified, 91
 - Vasicek and CIR models, 686–687
- Duration matching, 143, **798**
- Duration-based hedge ratio, 142
- Duration-based hedging strategies, 142–144
- DV01, **798**
- Dynamic hedging, 381, 383–386, 396, **798**
- Dynamic models, credit derivatives, 570
- Early exercise, 196, 225–227, 333, **798**
- Efficient market hypothesis, **798**
- Electricity derivatives, 1, 751–752
- Electronic trading, 2, 203, **798**
- Embedded option, 742–744, **798**
- Empirical research, **798**
- Employee stock option, 209, 316–317, 332–342, **798**
 - accounting for, 335–336
 - agency costs, 334, **791**
 - backdating, 341–342, **792**
 - binomial tree, 338
 - contractual arrangements, 332–333
 - early exercise of, 333
 - exercise multiple, 339–340
 - expected life, 337
 - repricing, 334
 - valuation of, 337–340
 - vesting, 333
 - Zions Bancorp, 340
- Energy derivatives, 750–752
 - hedge risks, 761
 - modeling energy prices, 752–757
- Enron’s bankruptcy, 537, 780
- Equilibrium model, interest rates, 683–689, **798**
 - application, 688–689
- Equity swap, 174, 740–741, **798**
- Equity tranche, 181–184, 560, **798**
- Equivalent martingale measure result, 630, 635–636
- Eurex, 22
- Eurocurrency, 76, **798**
- Eurodollar, 137, **798**
- Eurodollar futures, 137–142, 159, **798**
 - and LIBOR zero curve, 141–142
 - convexity adjustment, 140–141
 - vs. forward, 139–141
- Eurodollar futures options, 363
- Eurodollar interest rate, 137, **798**
- Euronext, 22, 199
- European option, 7, 194, **798**
 - binomial trees, 253–257
 - Black–Scholes–Merton model, on a non-dividend-paying stock, 313–316
 - Black–Scholes–Merton model, on a stock paying a known dividend yield, 350–352
 - delta, 382–383
 - dividend-paying stock, 229, 320–321
 - futures option compared to spot option, 361, 364–365
 - non-dividend-paying stock, 313–316
 - pricing using Black’s model, 371–372, 641
 - put–call parity, 221–224, 230, 236, 351, 409–410
 - risk-neutral valuation, 311–313
 - stock paying a known dividend yield, 350–352
- EWMA, *see* Exponentially weighted moving average.
- Excess cost layer, reinsurance, 759
- Excess-of-loss reinsurance contract, 759
- Exchange option, 586–587, 642–643, **798**
- Exchange rates
 - Black–Scholes–Merton and, 355–356
 - empirical data, 412–413
- Exchange-traded market, 2
 - difference between over-the-counter market and, 2–4
 - for options, 203
- Ex-dividend date, 229, 320, **798**
- Executive stock option, *see* Employee stock option
- Exercise boundary parameterization approach, 624–625

- Exercise limit, 203, **798**
- Exercise multiple, 339–340, **799**
- Exercise price, 7, **799**
- Exotic options, 210, 574–594, **799**
 - Asian options, 584–586, 611, **791**
 - barrier options, 579–581, 591–593, 615–618, **792**
 - basket options, 587, **792**
 - binary options, 581–582, **792**
 - chooser options, 578–579, **794**
 - cliquet options, 577
 - compound options, 577–578, 785
 - forward start options, 576–577, **800**
 - gap options, 575–576
 - lookback options, 582–584, 611–613, **803**
 - nonstandard American option, 575
 - options to exchange one asset for another, 586–587
 - packages, 574, **805**
 - Parisian options, 581
 - shout options, 584, **807**
- Exotics, *see* Exotic options.
- Expansion option, 770, 774–775
- Expectations theory, shape of zero curve, 93, **799**
- Expected life, 337
- Expected return, stock's, 289–290, 302–303
 - stock option price and, 257–258
- Expected shortfall, 472, **799**
- Expected spot price, 121
- Expected value of a variable, **799**
- Expiration date, 7, 194, **799**
 - effect on option price, 215
- Explicit finite difference method, 459–464, **799**
 - relation to trinomial tree approach, 453–454
- Exponential weighting, **799**
- Exponentially weighted moving average (EWMA), 500–502, **799**
 - application, 515
 - compared with GARCH, 502–504
 - estimating parameters, maximum likelihood methods, 504
- Exposure, credit 172, 532, **799**
- Extendable bond, **799**
- Extendable swap, 175, **799**
- Factor, 491–494, **799**
- Factor analysis, 490–494, **799**
- Factor loading, 491
- Factor score, 491
- FAS 123, 335, **799**
- FAS 133, 40, **799**
- FASB, *see* Financial Accounting Standards Board.
- FBI, 39
- Federal National Mortgage Association (FNMA), 728
- Federal Reserve Board, 38
- FICO, 186, **799**
- Fill-or-kill order, 38
- Financial Accounting Standards Board (FASB), 40, 335, **799**
- Financial intermediary, 152–153, **799**
- Finite difference method, 455–466, **799**
 - applications of, 465–466
 - explicit, 459–464
 - implicit, 456–459
 - other, 464–465
 - relation to trinomial tree approach, 453–454
- First notice day, 36
- First-to-default basket credit default swap, 558
- Fitch, 521
- Flat volatility, 656–657, **799**
- Flex option, 201, **799**
- Flexi cap, 724, **799**
- Flexible forwards, 575
- Flight to quality, 96, 786
- Floating lookback, 582
- Floor, 654–655, **800**. *See also* Interest rate caps and floors.
- Floor rate, **800**
- Floor–ceiling agreement, *see* Collar.
- Floorlet, interest rate, 654, **800**
- Fixed lookback, 583
- Force of interest, 79
- Foreign currency option, *see* Currency option.
- Forward band, 575
- Forward contract, 5–7, **800**
 - credit default swap, 557
 - delivery price, 6
 - delta, 398–399
 - futures vs., 7, 22, 41–42, 111–112
 - option vs., 8, 194
 - foreign exchange quotes, 5, 42
 - hedging, using, 10–11
 - valuing, 109–111, 312
 - VaR and, 484
- Forward difference approximation, 456
- Forward exchange rate, **800**
- Forward interest rate, 84–86, 162–164, 668–669, **800**
 - future interest rates vs., 139–140
- Forward price, 6, 104–111, **800**
 - for an investment asset that provides known cash income, 107–109
 - for an investment asset that provides known yield, 109

- for an investment asset that provides no income, 104–106
- Itô's lemma, applied to, 292
- relation to futures price, 111–112
- Forward rate, **800**
- Forward rate agreement (FRA), 86–89, 162–164 **800**
- Forward rate volatilities, 720–721
- Forward risk-neutral, 636–640, 669, 719
- Forward risk-neutral world, 636, **800**
- Forward start option, 576–577, **800**
- Forward swap, 174, 659, **800**
- Forward with optional exit, 575
- Front office, 784
- Front running, 39
- Full-service broker, 204
- Futures commission merchants (FCMs), 37, **800**
- Futures contract, 7, 22–43, **800**
 - asset underlying, 24–25
 - closing out positions, 23–24
 - commodities, 117–120
 - contract size, 25
 - currencies, 114–117
 - daily settlement, 27
 - delivery, 24, 36–37, 121
 - delivery month, 25
 - delta, 399
 - forward contracts vs., 7, 22, 41–42, 111–112
 - options vs., 7, 194
 - foreign exchange quotes, 42
 - index, 60–65, **802**
 - long position, 23
 - margins and, 27–30
 - marking to market, 27
 - price quotes, 26
 - risk and return, 122
 - short position, 23
 - specification of, 24–25
 - Treasury bond and Treasury note futures, 24–25, 132–136
- Futures, interest rate, 129–147
- Futures market, regulation of, 38
- Futures option, 199, 361–374, **800**
 - futures-style option, 372–373
 - interest rate futures option, 363–364
 - popularity of, 364
 - put–call parity, 365–366
 - spot options compared to, 361
 - valuation, using binomial trees, 272, 367–369, 435–437
 - valuation, using Black's model, 370–372
- Futures price, 23, 111–124, **800**
 - convergence to spot price, 26
 - cost of carry, 120
 - expected future spot prices and, 121–123
 - expected growth rate, 369–370
 - patterns of, 35
 - relationship to forward prices, 111–112
 - stock indices, 112–114
- Futures-style option, 372–373, **800**
- Gamma, 389, 398, **800**
 - cross gamma, 488
 - effect on VaR estimates, 486–488
 - estimating, using binomial tree, 434
 - interest rate derivatives, 664
 - relationship with delta and theta, 392–393
- Gamma-neutral portfolio, 392, **800**
- GAP management, 144
- Gap option, 575–576, **800**
 - application to insurance, 576
- GARCH model, 502–512, **800**
 - autocorrelation and, 508
 - compared with EWMA, 502–504
 - estimating parameters using maximum likelihood methods, 505–509
 - forecasting future volatility, 509–511
 - Ljung–Box statistic, 509
- Gaussian copula model, 538–540, 541, **800**
- Gaussian quadrature, 564, **800**
- Generalized Wiener process, 284–286, **800**
- Geometric average, **801**
- Geometric Brownian motion, 287, **801**
- Gibson Greetings, 781, 783, 789
- Ginnie Mae, *see* GNMA.
- Girsanov's theorem, 267, **801**
- GNMA, 180–181
- Good-till-canceled order, 38
- Google, 8–9
- Government National Mortgage Association (GNMA), 728
- Greek letters, Greeks, 377–404, **801**. *See also* Delta, Theta, Gamma, Vega, and Rho.
 - estimating, using binomial tree, 433–434
 - estimating, using finite difference method, 466
 - estimating, using Monte Carlo simulation, 451–452
 - interest rate derivatives, 664
 - Taylor series expansions and, 408
 - volatility smile and, 417–418
- Greenspan, Alan, 779
- Gross basis, 30
- Growth factor, 429
- Gucci Group, 202
- Haircut, 536, **801**
- Hammersmith and Fulham, 173, 781, 788
- Hazard rate, 522, **801**

- HDD, *see* Heating degree days.
- Heath, Jarrow, and Morton (HJM) model, 715–718
- Heating degree days (HDD), 758, **801**
- Hedge accounting, 39–41
- Hedge funds, 11, **801**
- Hedge ratio, **801**
- Hedge-and-forget, 47, 381
- Hedgers, 10–12, 789, **801**
- Hedging, 47–68, **801**
- and stock picking, 65
 - arguments for and against, 49–52
 - basic principles, 47–49
 - basis risk, 52–56
 - competitors and, 50–51
 - cross, 56–60
 - delta hedging, 264–265, 380–387
 - duration-based hedging strategies, 142–144
 - equity portfolio, 62–65
 - exotic options, 591–593
 - gamma, 389–392
 - gold mining companies and, 52
 - hedge and forget, 47, 381
 - hedge effectiveness, 58
 - hedge ratio, 56
 - in practice, 396–397
 - interest rate derivatives, 663–664, 709–710
 - long hedge, 48–49
 - Metallgesellschaft (MG) and, 67
 - naked and covered position, 378
 - perfect hedge, 47
 - performance measure, 360, 385
 - rho, 395–396
 - rolling forward, 65–67
 - shareholders and, 50
 - stack and roll, 65–67
 - static options replication, 591–593, **807**
 - stop-loss strategy, 378–380
 - short hedge, 48
 - theta, 387–389
 - using index futures, 60–65
 - vega, 393–395
- Heston model, 606
- Heterogeneous model, 569
- High-frequency trading, 2
- Historical volatility, 304–306, **801**
- Historical simulation, value at risk, 474–478, **801**
- compared with model-building approach, 489–490
- History-dependent derivative, 611
- HJM, *see* Heath, Jarrow, and Morton model.
- Ho–Lee model, 690–691
- Holiday calendar, 155, **801**
- Homogeneous model, 569
- Hopscotch method, 464
- Hull–White employee stock option model, 339–340
- Hull–White interest rate model
- one-factor, 691–693
 - two-factor, 694
- Hull–White, dealing with path dependence, 611
- Hull–White stochastic volatility model, 606
- Hunt brothers, 39
- Hurricane Andrew, 759
- IASB, *see* International Accounting Standard Board, 40
- IAS 2, 335
- IAS 39, 40
- IMM dates, **801**
- Implicit finite difference method, 456–459, **801**
- Implied copula, 570
- Implied correlation, 566–568, **801**
- Implied distribution, **801**
- currency options, 411–413
 - determining, 424–426
 - stock options, 414–415
- Implied dividend yield, 354–355
- Implied tree model, **607, 801**
- Implied volatility, 318–320, **801**
- Implied volatility function (IVF) model, 607–608, **801**
- Importance sampling, variance reduction procedure, 453–454
- Incentives, 788
- Inception profit, 784–785, **801**
- Index amortizing swap, 744, **801**
- Index arbitrage, 113–114, **802**
- Index currency option note (ICON), 19
- Index futures, 60–61, **802**
- changing portfolio beta, 64–65
 - hedging, using index futures, 60–65
 - portfolio insurance, 401–402
 - pricing, 112–114
 - quotations, 61
- Index option, 199, 345–347, **802**
- implied dividend yield, 354–355
 - portfolio insurance, 401–402
 - valuation, binomial tree, 270, 435–437
 - valuation, Black–Scholes–Merton, 352–355
- Indexed principal swap, 744, **802**
- Initial margin, 27, 205, **802**
- Inner barrier, 615–616
- Instantaneous forward rate, **802**
- Instantaneous short rate, 682. *See also* short rate.

- Insurance derivatives, 1, 759
 - pricing, 760–761
- IntercontinentalExchange (ICE), 24, 750
- Interest only (IO), 730, **802**
- Interest rates, 75–97
 - continuous compounding, 79–80
 - convexity adjustment, 668–672
 - day count conventions, 129–130, 154–155
 - equilibrium models, 683–689
 - forward, 84–86
 - forward rate agreements (FRA), 86–89
 - short-rate models, 682–710
 - spot, 80
 - term structure theories, 93–96
 - types of, 75–77
 - zero-coupon yield curve, 82–84
- Interest rate caps and floors, 653–659, **802**
 - cap as portfolio of bond options, 654
 - cap as portfolio of interest rate options, 653
 - floors and collars, 654
 - impact of day count conventions, 658–659
 - put–call parity, 655
 - spot volatilities vs. flat volatilities, 656–657
 - theoretical justification of Black’s model, 657–658
 - valuation of cap and floors, 654–656
- Interest rate collar, 654, **802**
- Interest rate derivatives, 648–665, **802**
 - American bond options, 648–652, 694–695, 707–708
 - calibration, 707–709
 - hedging, 663–664, 709–710
 - interest rate caps and floors, 653–659
 - mortgage-backed securities, 728–730
 - no-arbitrage models, 689–694
 - spot volatilities vs. flat volatilities, 656–657
 - swap option, 659–663, 659–663, 724–728, 727–728
 - tree-building procedure, 698–709
 - trees for, 696–709
 - volatility structures, 695–696, 707–709, 720–721, 726–728
- Interest rate floor, 654, **802**. *See also* Interest rate caps and floors.
- Interest rate futures, 129–147
 - Eurodollar futures, 137–142
 - relation to forward interest rate, 139–140
 - Treasury bond futures, 132–136
- Interest rate futures option, 363–364
- Interest rate model
 - Black, 641, 648–663
 - Black–Karasinski, 693–694
 - Brennan–Schwartz
 - calibration, 707–709
 - Cox–Ingersoll–Ross, 685–689
 - Heath–Jarrow–Morton, 715–718
 - Ho–Lee, 690–691
 - Hull–White (one-factor), 691–693
 - Hull–White (two-factor), 694
 - LIBOR market model, 718–728
 - Longstaff–Schwartz, 686
 - Rendleman–Bartter, 683–684
 - Vasicek, 684–689
- Interest rate options, 353–364, 648–731, **802**
- Interest rate parity, 115
- Interest rate swap, 148–165, 733–740, **802**
 - comparative-advantage argument, 156–158
 - confirmation, 155
 - day count conventions, 154–155
 - financial intermediary, role of, 152–153
 - mechanics of, 148–154
 - plain vanilla interest rate swap, 148
 - to transform a liability, 151
 - to transform an asset, 151–152
 - valuation, 160–164
 - VaR and, 484
- Interest rate trees, 696–709
- Internal controls, 790
- International Accounting Standards Board, 40, 335
- International Securities Exchange, 199
- International Swaps and Derivatives Association (ISDA), 155, 550, **802**
 - Master Agreement, 155
- In-the-money option, 201, **802**
 - delta, 383
 - gamma, 392
 - theta, 388
- Intrinsic value, 201, **802**
- Inverted market, 35, **802**
- Investment asset, 101, **802**
 - forward price for an, 104–111
 - market price of risk, 632–633
- Investment grade, 521
- ISDA, *see* International Swaps and Derivatives Association.
- Itô process, 286, **802**
- Itô’s lemma, 291–293, 634, 636, **802**
 - extensions, 298
 - proof, 297–298
- iTraxx Europe, 555, 561, **802**
- IVF, *see* Implied volatility function model.
- J. P. Morgan, 501
- Jump–diffusion model, 600–603, **802**
- Jump process, 601–605, **803**
- Kidder Peabody, 106, 779, 783, 784
- Knock-in and knock-out options, 579
- Kurtosis, 411, **803**

- Last notice day, 36
- Last trading day, 36
- LEAPS (Long-term equity anticipation securities), 200, 345, **803**
- Least-squares approach, 621–624
- Lehman bankruptcy, 3, 550, 787
- Levenberg–Marquardt procedure, 708
- Levy processes, 600
- LIBID, 76, **803**
- LIBOR, 76, **803**
 - OIS rate vs., 188
 - swap rate vs., 159
- LIBOR curve, 141–142, 159–160, **803**
- LIBOR market model, interest rates derivatives, 718–728
- LIBOR-in-arrears swap, 174, 737–738, **803**
- Limit move, 26, **803**
 - limit down, 26
 - limit up, 26
- Limit order, 37, **803**
- Linear model, value at risk, 481–486
- Liquidity preference theory, 94–95, **803**
- Liquidity premium, **803**
- Liquidity risk, 95–96, 785–787, **803**
- Lloyds syndicate, 759
- LMM, *see* LIBOR market model.
- Locals, 37, **803**
- Lock-out period, 575, 649
- Lognormal distribution, 293, 301, **803**
- Lognormal property, 292–293, 300–301
- London interbank bid rate (LIBID), 76, **803**
- London interbank offered rate (LIBOR), 76, **803**
 - zero curve, 141–142, 159–160, **803**.
- London Stock Exchange, 15
- Long hedge, 48–49, **803**
- Long position 5, **803**
 - option, 9
- Longstaff–Schwartz interest rate model, 686
- Longstaff–Schwartz simulation model, 621–624
- Long-Term Capital Management (LTCM), 31, 780, 785–786
- Long-term equity anticipation securities (LEAPS), 200, **803**
- Lookback option, 582–584, 611–613, **803**
 - fixed, 583
 - floating, 582
- Low-discrepancy sequence, 455, 785
- Maintenance margin, 28, 205, **803**
- Margin, 27–30, 205–206, **803**
 - buying on margin, 205
 - clearing margin, 30
 - futures contracts, 27–30
 - gross margining, 30
 - initial margin, 27, 205
 - maintenance margin, 28, 205
 - margin call, 28, 206, **803**
 - margin requirements, 205
 - net margining, 30
 - stock options, 205
 - variation margin, 28
- Margin account, 13, 27, 103, 206
- Margrabe model, 586
- Market maker, 153–154, 203, **803**
- Market model, **803**
- Market order, 37
- Market price of risk, 631–634, 768–769, **804**
 - several state variables, 634–635, 640–641
- Market segmentation theory, 93–94, **804**
- Market-if-touched (MIT) order, 37
- Market-not-held order, 38
- Market transparency, 787–788
- Marking to market, **804**
- Marking to model, 784
- Markov process/property, 280–281, **804**
- Martingale, 630, 635–636, **804**
 - equivalent martingale measure result, 630, 635–636
- Maturity date, 7, 194, **804**
- Maximum likelihood method, 504–508, 688, **804**
- Mean reversion, 503, 684, **804**
 - commodity, 749, 753–755
- Measure, 630, **804**
- Merton’s model (debt as option on assets of firm), 530–531
- Merton’s mixed jump–diffusion model, 601–603
- Metallgesellschaft (MG), 67, 781
- Metals, 749–750
- Mezzanine tranche, 181–183, 560, **804**
- Mid-curve Eurodollar futures option, 362
- Middle office, 784
- Midland Bank, 780
- Minimum variance hedge ratio, 57–58
- Min–max, 575
- Mixed jump–diffusion model, 601–603
- Model-building approach, value at risk, 478–481
 - compared with historical simulation, 489–490
- Modified duration, 91, **804**
- Moment matching, variance reduction procedure, 454–455
- Moments,
 - Asian options, 585
 - basket options, 588
- Money market account, 637, **804**

- Monte Carlo simulation, 288–289, 380, 427, 446–455, 488–489, 611, 621–625, **804**
 American options and, 621–625
 Black–Scholes–Merton model and, 448–449
 calculating π with, 446
 exercise boundary parameterization approach, 624–625
 generating random samples, 450–451
 Greek letters and, 451–452
 least-squares approach, 621–624
 LIBOR market model of short rates, 721–723, 726
 number of trials, 450–451
 partial simulation approach (VaR), 489
 value-at-risk measure, 488–489
 valuing derivatives on more than one market variable, 448–449
 valuing mortgage-backed securities, 729
 valuing new business, 769, 776
- Mortgage-backed security (MBS), 180–181, 728–730, **804**
 collateralized mortgage obligations, 729
 option-adjusted spread (OAS), 729–730
 stripped mortgage-backed security, 730
 valuing mortgage-backed securities, 729
- Moody's, 521
- Mutual fund's return, 304
- Naked option, 205–206
- Naked position, 378, **804**
- Nasdaq 100 index (NDX), 61, 199
 index options, 199
 Mini Nasdaq 100 index futures, 61
- Nasdaq OMX, 199, 347
- National Futures Association (NFA), 38
- Natural gas derivatives, 750–751
- Net basis, 30
- Net interest income, 94
- Netting, 534–535, **804**
- Neutral calendar spread, 245
- New York Federal Reserve, 780
- New York Mercantile Exchange (NYMEX), 750
- New York Stock Exchange, *see* NYSE Euronext.
- Newton–Raphson method, 81, 318, 705, **804**
- Nikkei futures, 113, 674–675, 780
- NINJA, 186, **804**
- No-arbitrage assumption, **804**
- No-arbitrage interest rate model, 689–694, **804**
 Black–Karasinski model, 693–694
 Heath–Jarrow–Morton (HJM) model, 715–718
 Ho–Lee model, 690–691
 Hull–White (one-factor) model, 691–693
 Hull–White (two-factor) model, 694
 LIBOR Market Model (LMM), 718–728
- Noncentral chi-squared distribution, 601
- Nonstandard American options, 575
- Nonstationary volatility structure, 709, **804**
- Nonsystematic risk, 73, 122, 635, **804**
- Normal backwardation, 123, **805**
- Normal distribution, 281, **805**
- Normal market, 35, **805**
- Northern Rock, 96, 787
- Notice of intention to deliver, 24
- Notional principal, 150, 548, **805**
- Numeraire, 636–640, **805**
 annuity factor as the numeraire, 639–640
 impact of a change in numeraire, 643–644
 interest rates when a bond price is the numeraire, 638–639
 money market account as the numeraire, 637
 numeraire ratio, 644
 zero-coupon bond price as the numeraire, 637–638
- Numerical procedure, 427–466, **805**
- NYSE Euronext, 22, 199
- OCC, *see* Options Clearing Corporation (OCC).
- October 19, 1987, 30, 114, 403, 416, 490, 786
- Offer, 3, 203, **805**
- Offsetting orders, 204
- Off-the-run bond, 786
- OIS, *see* Overnight-indexed swap.
- On-the-run bond, 786
- Open interest, 35, 204, **805**
- Open order, 38
- Open outcry trading system, 2, **805**
- Option, 7–9, **805**
 class, 201, **805**
 credit default swap, 558
 difference between futures (or forward) contracts and, 8, 194
 exercise limits, 203
 exercising, 207
 exotic, 210, 574–594
 fence, 575
 foreign currency, *see* Currency options.
 Futures, *see* Futures options.
 hedging, using, 12
 index, *see* Index options.
 intrinsic value, 201
 on futures, *see* Futures options.
 position limits, 203
 positions, 196–198
 regulation of, 207
 series, 201, **805**

- stock, *see* Stock options.
- swap, *see* Swaptions.
- taxation, 207–209
- time value, 201
- trading, 203–204
- two correlated assets, 618–621
- types of, 7, 194
- valuation, *see* Stock option valuation.
- Option-adjusted spread (OAS), 729–730, **805**
- Options Clearing Corporation (OCC), 206, **805**
- Options in an investment opportunity, 769–776
- Options involving several assets, 587–588
- Options on bonds, 694–695, 707–708
- Options to defer, 771
- Options to exchange one asset for another, 586–587
 - forward risk-neutral valuation, 642–643
- Options to extend, 771
- Orange County, 87, 779, 781, 782, 788
- Order book official, **805**
- Order, types of, 37–38
- OTC market, *see* Over-the-counter market.
- Outer barrier, 615–616
- Out-of-the-money option, 201, **805**
 - delta, 383
 - gamma, 392
 - theta, 388
- Outside model hedging, 710
- Overnight repo, 76
- Overnight-indexed swap, 77, 164–165, **805**
- Over-collateralization, 183
- Over-the-counter market, 3–4, 30–31 **805**
 - options, 210
- Packages, 574, **805**
- Par value, **805**
- Par yield, 81–82, **805**
- Parallel shift, **805**
- Parisian options, 581, **805**
- Partial simulation approach, Monte Carlo simulation, 489
- Pass-throughs, 729
- Path-dependent derivative, 611–615, **805**
- Payoff, **806**
- PD, **806**
- Philadelphia Stock Exchange (PHLX), 199
- Physical default probabilities, 528
- Plain vanilla product, 148, 574, **806**
- P-measure, 267, **806**
- Poisson process, 601, **806**
- Portfolio immunization, 143, **806**
- Portfolio insurance, 345–347, 400–402, 786, **806**
 - stock market volatility and, 402
- Position limit, 26, 203, **806**
- Positive-semidefinite matrix, 514
- Position traders, 37
- Premium, **806**
- Prepayment function, 728, **806**
- Prices
 - Price sensitivity hedge ratio, 142
 - opening price, 33
 - settlement price, 35, **807**
- Principal, **806**
- Principal components analysis, 490–494, **806**
- Principal only (PO), 730, **806**
- Principal protected note, 234–236, **806**
- Probability measure, 634
- Probability of default, 522–531
 - comparison of default probabilities estimates, 526–530
 - estimating, using bond prices, 524–526
 - estimating, using equity prices, 530–531
 - historical default probabilities, 522–523
 - implied from asset swaps, 526
 - implied from bond data, 524–526
 - implied from credit default swaps, 554
 - risk-neutral vs. real-world, 528–529
- Procter and Gamble, 744–745, 781, 783, 789
- Profit center, 789
- Program trading, 114, **806**
- Protective put, 236, **806**
- Pull-to-par, **806**
- Pure jump model, 600
- Put–call parity, 221–225, 230, 236, 351, 354, 365–366, 409–410, **806**
 - caps and floors, 655
- Put option, 7, 194–196, **806**
- Puttable bond, 649, **806**
- Puttable swap, 175, **806**
- Q-measure, 267, **806**
- Quadratic model, value at risk, 486–488
- Quadratic resampling, variance reduction procedure, 454
- Quanto, 174, 674–677, **806**
- Quasi-random sequence, variance reduction procedure, 455, **806**
- Quotations
 - commodity futures, 34
 - currency futures, 42, 117
 - foreign exchange rate, 42
 - interest rate futures, 133–134
 - stock index futures, 61
 - Treasury bills, 131
 - Treasury bond and note futures, 133
 - Treasury bonds, 131
 - US dollar swaps, 154
 - USD–GBP exchange rate, 5

- Quoted price, bond and Treasury bill, 131, 650, 707. *See also* Clean price.
- Rainbow option, 587, **806**
- Random factor loadings model, 569
- Random walk, 253
- Range forward contract, 348–350, 575, **807**
- Ratchet cap, 723, **807**
- Rating agencies, 76, 521
- Real options, 1, 765–776, **807**
- Real world
 - commodities, 748
 - interest rates, 682, 688–689
- Rebalancing, 308, 381, **807**
- Recovery rate, 523–524, 554–555, **807**
- Reference entity, 548, **807**
- Regulatory arbitrage, 189
- Reinsurance, against catastrophic risks (CAT reinsurance), 759
- Rendleman–Bartter interest rate model, 683–684
- Repo, 76, **807**
- Repo rate, 76, **807**
 - overnight repo, 76
 - term repo, 76
- Repricing (employee stock options), 334
- Repurchase agreement, 76
- Reset date, **807**
- Retractable bond, 649
- Reverse calendar spreads, 245
- Reversion level, **807**
- Rho, 395–396, 398, **807**
 - estimating, using binomial tree, 434
- Rights issue, **807**
- Right-way risk, 534
- Risk, *see also* Value at risk,
 - back testing, 490
 - basis, 52–56
 - credit, 171–172, 521–543
 - foreign exchange, 10–11
 - nonsystematic, 73, 122, 635, **805**
 - stress test, 490, 783
 - systematic, 73, 122, 635, **808**,
- Risk and return, relationship between for futures, 122
- Risk-free interest rate, 76–77, 525–526, **807**
 - effect on option price, 217
- Risk limits, 781–782
- RiskMetrics, 501
- Risk-neutral valuation, 257–259, 311–313, 352, 428, 528, 748, 766–768, **807**
- Risk-neutral world, 257, 581, **807**. *See also* Forward risk-neutral.
 - interest rates, 682, 688–689
 - real world vs., 259
 - rolling forward risk-neutral, 719
 - traditional, 630, 633
- Roll back, **807**
- Rolling forward risk-neutral, 719
- Russell 1000 index, 61
- Scalper, 37, **807**
- Scenario analysis, 397, 783, **807**
- Seasonality, 749, 755–757
- Securities and Exchange Commission (SEC), 38, 207, 340, **807**
- Securitization, 180–184, **807**
- Segmentation theory, shape of zero curve, 93–94
- Settlement price, 35, **807**
- Sharpe ratio, 632, **807**
- Shell, 781
- Short hedge, 48, **807**
- Short position, 5, 9, **807**
- Short rate, 682, **807**
 - calibration, 707–709
 - equilibrium models, 683–689
 - general tree-building procedure, 698–709
 - hedging using one-factor models, 709–710
 - interest rate trees, 696–709
 - no-arbitrage models, 689–694
 - options on bonds, 694–695
 - volatility structures, 695–696
- Short selling, 102–103, **808**
- Shorting, 102–103
- Short-term risk-free rate, *see* Short rate.
- Short-term funding, 787
- Shout option, 584, **808**
- Siegel's paradox, 676
- Simulation, **808**
- Single-tranche trading, 560–561
- Société Générale, 17, 779, 780, 784
- Special purpose vehicle (SPV), 189
- Specialist, **808**
- Speculation
 - using futures, 13
 - using options, 14
- Speculators, 13–15, 789, **808**
- Spot
 - contract, 5
 - convergence of futures price to spot price, 26
 - forward prices and spot prices, 6, 104–111
 - futures option compared to spot option, 364–365
 - futures prices and expected future spot prices, 121–123
 - interest rate, 80, **808**
 - price, **808**
 - traders, 6
 - volatility, 656–657, **808**

- Spread option, **808**
- Spread trading strategy, 238–246
- Spread transaction, 29, **808**
- SPV, *see* Special purpose vehicle.
- Stack and roll, 65–67, **808**
- Standard and Poor's (S&P), 521
- Standard and Poor's (S&P) Index, 61
 - 100 Index (OEX and XEO), 199, 345
 - 500 Index (SPX), 61, 199, 345
 - 500 Index futures, 36
 - Mini S&P 500 futures, 61
 - options, 199, 345
- Static hedge, 381, 591–593, **808**
- Static options replication, 591–593, **808**
- Step-up swap, 174, 733, **808**
- Sticky cap, 723, **808**
- Stochastic process, 280–294, **808**
 - correlated, 290–291
- Stochastic variable, **808**
- Stochastic volatility models, 605–607
- Stock dividend, 202–203, **808**
- Stock index/indices, 61, **808**
- Stock index futures, *see* Index futures.
- Stock index option. *See* Index option.
- Stock option, 199–203, **808**. *See also* Black–Scholes–Merton model.
 - commissions, 204–205
 - dividend and stock split, 201–203
 - employee, 209, 332–342
 - expiration dates, 194, 200
 - flex option, 201, 345
 - long-term equity anticipation securities (LEAPS), 200, 345
 - margins, 205
 - naked, 205–206
 - position and exercise limits, 203
 - regulations of, 207
 - specification of, 199–203
 - strike price, 200
 - taxation, 207–209
 - terminology, 201
 - trading, 203–204
- Stock option valuation,
 - American options on dividend-paying stock, 321–323
 - American options on non-dividend-paying stock, 227–229, 429–433
 - assumptions, 218
 - binomial tree, 253–257, 427–434
 - Black's approximation, 322–323
 - bounds for dividend-paying stocks, 229
 - bounds for non-dividend-paying stocks, 218–221
 - dividend yield, 350–352
 - dividends, 201–203, 217–218, 320–323
 - employee stock options, 337–340
 - European options on a dividend-paying stock, 320–321, 350–352
 - European options on a non-dividend-paying stock, 221–222, 313–316
 - European options on stocks paying known dividend yields, 350–352
 - factors affecting prices, 214–218
 - implied distribution and lognormal distribution, 414–415
 - put–call parity, 221–225, 230, 409–410
 - risk-neutral valuation, 352
 - single large jump in asset anticipated, 419–420
 - stock's expected return and, 257–258
 - volatility smile (skew), 414–415
- Stock prices
 - expected return, 302–303
 - lognormal property, 292–293, 300–301
 - process for, 286–290
 - rate of return, distribution of, 301–302
 - volatility, 303, 414–415
- Stock split, 201–202, **808**
- Stock's expected return, 302–303
 - irrelevance of, 257–258, 311–312
- Stocks-to-use ratio, 749
- Stop-and-limit order, 37
- Stop-limit order, 37
- Stop-loss order, 37
- Stop-loss strategy, 378–380
- Stop order, 37
- Storage cost, 118, **808**
- Straddle, 246–247, **808**
- Strangle, 247, **808**
- Strap, 247, **808**
- Stratified sampling, variance reduction procedure, 454
- Strengthening of the basis, 53
- Stressed VaR, 490, **809**
- Stress testing, 191, 490, 783, **809**
- Strike price, 7, 194, **809**
- Strip, 82, 247, **809**
- Strip bonds, 82, 106, **809**
- Stripped mortgage-backed securities, 730
- Subprime mortgage, 780, 787, **809**
- Sumitomo, 781
- Swaps, 148–175, **809**
 - accrual, 174, 742–743
 - amortizing, 174, 733
 - basis, 734
 - cancelable, 743
 - cancelable compounding, 743
 - clearing house, 172
 - commodity, 175, 744

- comparative-advantage argument, 156–158, 166–168
- compounding, 174, 735–736
- confirmations, 155
- constant maturity (CMS), 174, 738–739
- constant maturity Treasury (CMT), 174, 739
- credit risk and, 171–172
- cross currency, 174, 737
- currency, 165–171, 736–737, **797**
- deferred, 659
- differential (diff swap), 740
- embedded options, 742–744
- equity, 174, 740–741
- forward swaps, 174, 659
- index amortizing rate, 744
- indexed principal, 744
- interest rate, *see* Interest rate swaps.
- LIBOR-in-arrears, 174, 737–738
- step-up, 174, 733
- total return, 558–559, **810**
- variance, 588–591
- volatility, 175, 588–591
- Swap option, *see* Swaption.
- Swap rate, 153, 659–660, **809**
 - convexity adjustment, 671–672
 - LIBOR rate vs., 159
- Swap zero curve, 159–160. *See also* LIBOR zero curve.
- Swaption, 175, 659–663, **809**
 - Bermudan swaptions, 727–728
 - European swaption, 659–663, 724–728
 - implied volatilities, 661–662
 - relation to bond options, 660
 - theoretical justification for Black's model, 662
- Swing option, electricity and natural gas market, 752, **809**
- Synthetic CDO, 560–565, **809**
 - valuation, 562–565
- Synthetic option, 400–402, **809**
- Systemic risk, 32–33, **809**
- Systematic risk, 73, 122, 529, 635, **809**
- Tailing the hedge, 60, **809**
- Tail loss, 472, **809**
- Take-and-pay option, electricity and natural gas market, 752, **809**
- Tax, 40–41
 - planning strategy, 209
- Taxpayer Relief Act of 1997, 208
- Teaser rate, 185
- TED spread, 188, **809**
- Tenor, 653
- Term repo, 76
- Term structure model of interest rates, 682, **809**
- Term structure theories, shape of zero curve, 93–96
- Terminal value, **809**
- Theta, 387–389, 398, **809**
 - estimating, using binomial tree, 434
 - relationship with delta and gamma, 392–393
- Time decay, 387, **809**
- Time value, 201, **809**
- Time-dependent parameters, 445–446
- Time-of-day order, 38
- Timing adjustment, 672–674, **809**
 - accrual swap, 742
 - constant maturity swap (CMS), 739
- To-arrive contract, 2
- Tokyo International Financial Futures Exchange (TIFFE), 22
- Top straddles, 247
- Top vertical combination, 248
- Total return swap, 558–559, **810**
- Tradeable derivatives, prices of, 311
- Traders, types of, 9–16, 37–38
- Trading days vs. calendar days, 306–307
- Trading strategies, involving options, 234–250
 - combinations, 246–248
 - for single option and stock, 236–237
 - spreads, 238–246
- Trading volume, 35
- Traditional risk-neutral world, 630, 633, 675–677
- Tranche, 181–184, 560, **810**
- Transaction cost, 387, **810**
- Treasury bill, **810**
- Treasury bond, **810**
- Treasury bond futures, 132–136, **810**
 - cheapest-to-deliver bond, 134–135
 - conversion factors, 133–134
 - quotations, 133
 - wild card play, 135
- Treasury bond futures option, 363–364
- Treasury note, **810**
- Treasury note futures, 132–136, **810**
- Treasury rate, 75–76
 - zero rate, 80, 82–84
- Tree, **810**. *See also* Binomial tree and Trinomial tree.
- Trinomial tree, 444–445, 696–709, **810**
 - analytic results, 705–706
 - calculation, 696–697
 - calibration, 707–709
 - commodity, 752–757
 - construction for Black-Karasinski, 703–705
 - construction for Hull-White, 698–703

- nonstandard branching, 697–698
- relation to finite difference method, 453–454
- valuation of American bond option, 707–708
- Triple witching hour, **810**
- Unconditional default probability, 522
- Underlying variable, **810**
- Unsystematic risk, 122, 635, **810**
- Up-and-in calls, 580
- Up-and-in option, 580, **810**
- Up-and-in puts, 580
- Up-and-out calls, 580
- Up-and-out option, 580, **810**
- Up-and-out puts, 580
- Uptick, 103, **810**
- US Department of Agriculture, 748
- US Department of Energy, 758
- US dollar index, 61
- US Treasury Department, 38
- Valuation of a business, 769
- Value at risk (VaR), 471–494, **810**
 - bank regulators use, 472
 - cash-flow mapping and, 483–484
 - comparison of approaches, 489–490
 - confidence level, 471
 - C-VaR, 472
 - diversification benefits and, 480–481
 - expected shortfall, 472, **799**
 - historical simulation, 474–478
 - interest rates and, 483–484
 - linear model, 481–486
 - model-building approach, 478–481
 - Monte Carlo simulation, 488–489
 - principal components analysis, 490–494
 - quadratic model, 486–488
 - RiskMetrics and, 501
 - time horizon, 473–474
 - volatilities and, 478–479
- Variance–covariance matrix, 483, 514, **810**
- Variance–gamma model, 603–605, **810**
- Variance rate, 498, **810**
 - estimating using maximum likelihood methods, 504–505
- Variance reduction procedures, 452–455, **810**
 - antithetic variable technique, 453
 - control variate technique, 453
 - importance sampling, 453–454
 - moment matching, 454–455
 - quadratic resampling, 454
 - quasi-random sequences, 455
 - stratified sampling, 454
- Variance swap, 588–591, **810**
- Variance targeting, 506
- Variation margin, 28, **810**
- Vasicek interest rate model, 684–689
 - compared with Cox, Ingersoll, Ross, 687
- Vasicek credit model, 540
- Vega, 393–395, 398, **810**
 - estimating, using binomial tree, 434
 - interest rate derivatives, 664
- Vega neutral portfolio, 393–394, **811**
- Vesting period, 333, **811**
- VIX index, 319–320, 590–591, **811**
- Volatility, interest rate derivatives, **811**
 - flat volatility, 656–657
 - forward rate volatility, 720–721
 - spot volatility, 656–657
 - structures, 695–696
 - volatility skews, 727, **811**
- Volatility, stock price,
 - Black–Scholes–Merton model and, 303, 409–421
 - causes of, 307
 - defined, 290, 303
 - effect on option price, 215
 - estimating from historical data, 304–306, 498–512. *See also* EWMA, GARCH.
 - forecast future volatility, 509–511
 - implied, 318–320. *See also* Volatility smile.
 - matching volatility with u and d , 265–267
 - portfolio insurance and stock market volatility, 402
 - swap, 175, 588–591, **811**
 - term structure, volatility of stock return, 416–417, 510–511, **811**
 - volatility surface, 416–417
 - vega, 393–395, **810**
- Volatility skew, 414, 727, **811**
- Volatility smile, 409–421, **811**
 - determining the implied asset price distribution, 424–426
 - equity options, 414–415
 - foreign currency options, 411–413
 - Greek letters and, 417–418
- Volatility surface, 416–417, **811**
- Volatility swap, 175, 588–591, **811**
- Volatility term structure, 416–417, 510–511, **811**
- Warrant, 209, 316–318, **811**
- Wash sale rule, 208
- Waterfall, 182, 559, **811**
- Weakening of the basis, 53
- Weather derivatives, 1, 758, **811**
 - pricing of, 760–761

Weather Risk Management Association
(WRMA), 758

Wiener process, 282–284, **811**

Wild card play, 135, **811**

Writing a covered call, 236

Writing an option, 9, **811**

Wrong-way risk, 534

Yield, **811**

Yield curve, **811**

Orange County and, 87

Yield volatilities, 652

Zero curve, 82–84, **811**

buckets and, 144

expectations theory, 93

liquidity preference theory, 94–95

market segmentation theory, 93–94

shape of zero curve, theories, 93–96

Zero rate, 80, **811**

Zero-cost collar, 575

Zero-coupon bond, **811**

Zero-coupon interest rate, 80, **811**

Zero-coupon yield curve, *see* Zero curve.

Zions Bancorp, 340