This looks like a complex formula, but it is actually just the formula for the forward rate using the LIBOR term structure. A typical forward rate calculation would have 1 + the longer term rate raised to a power in the numerator and 1 + the shorter term rate raised to a power in the denominator. This is more or less what we see here, but there are no exponents because compounding and discounting in the Eurodollar market is done using rate × days/360, rather than raising 1 + the rate to a power. Also, in the above formula subtracting 1 and multiplying by 360/m annualizes the rate.

Table 13.1 illustrates how this formula is used to solve for the rate on an FRA. We use the same example given earlier. A firm enters into an FRA that expires in 30 days. The underlying is 90-day LIBOR. The term structure of interest rates is 11 percent for 30 days and 10.32 percent for 120 days.

Recall that in examining swaps, the value of the contract is zero at the start. Later during its life—after the swap was initiated but before it expires—we needed to determine the value of the swap. We need to do the same for the FRA. To do so, let us position ourselves at a future day g, which is after day 0 but before the expiration, day h. We hold an FRA set at the rate F that will pay off in h − g days. The relevant term structure is LIBOR for h − g days and h + m − g days. Recall the transactions that we set up on day 0 to replicate the payoff of the FRA. The value of that combination of Eurodollar time deposits on day g will give us the value of the FRA on day g. Remember that the combination consists of a time deposit that we owe paying 1 + F(m/360) on day h + m and one that we hold paying $1 on day h. The value of this combination on day g is the value of the FRA, which we denote as VFRA,

[redacted]

Table 13.1 Solving for the Rate on a Forward Rate Agreement

To solve for the rate on a 30-day FRA in which the underlying is 90-day LIBOR, we have h = 30 and m = 90. Thus, we need $L_0(30)$ and $L_0(120)$, the 30- and 120-day LIBOR spot rates, which are

Term	Rate
30 days	11.00%
120 days	10.32%

The forward rate is found as follows:

$$F = \left(\frac{1 + 0.1032\left(\frac{120}{360}\right)}{1 + 0.11\left(\frac{30}{360}\right)} - 1 \right)\left(\frac{360}{90}\right) = 0.10.$$

The first term in the large parentheses is the value of the Eurodollar time deposit that we hold. Its value is based on discounting the payoff of $1 at the rate for h − g days. The second term in large parentheses is the value of the time deposit we owe, in which we pay 1 + F(m/360) in h + m − g days, discounted at the rate for h + m − g days.

An alternative expression for the value of an FRA on day g is (note we introduce a subscript to denote when the forward rate is set)

$$\text{VFRA} = \frac{(F_g - F_0)\left(\frac{m}{360}\right)}{1 + L_g(h + m - g)\left(\frac{h + m - g}{360}\right)}.$$

The value of the FRA is the present value of the change in the forward rate.

Table 13.2 illustrates this result. We see that changes in market conditions over this 20-day period of time have resulted in an increase in the value of the FRA from zero to over $3,000.

Applications of FRAs

In the example we have been working with, the buyer of the FRA could well have been a pure speculator, using the instrument to profit from an expectation of higher interest rates. In most cases, however, users of FRAs are borrowers and lenders who want protection against interest rate changes. FRAs are ideal for parties who anticipate the need to borrow money at a future date. The most appropriate type of loan for an FRA is a loan equivalent to a zero coupon bond. The borrower receives the proceeds from the loan and makes a single payment at maturity of principal plus interest. Of course, for the best match, the interest on the loan should be calculated in the same manner in which it is calculated on the FRA.

Table 13.2 Valuing a Forward Rate Agreement During Its Life

Consider the FRA we priced in Table 13.1. It is now 20 days into its life, with 10 days remaining. Thus, g = 20, h = 30, and h + m = 120. We shall need the term structure for h − g = 10 days and h + m − g = 100 days. These rates are

Term	Rate
10 days	10.55%
100 days	10.15%

The value of the FRA is found as follows:

$$20{,}000{,}000\left[\left(\frac{1}{1+0.1055\left(\frac{10}{360}\right)}\right)-\left(\frac{1+0.10\left(\frac{90}{360}\right)}{1+0.1015\left(\frac{100}{360}\right)}\right)\right]=\$3{,}697$$

Continuing with our example, consider a firm that needs to borrow $20 million in 30 days. The rate will be set at LIBOR plus 100 basis points. The firm is concerned with the risk associated with LIBOR over this 30-day period, so it feels that the purchase of an FRA would be a good way to protect against a possible increase in LIBOR. Of course, the FRA would also prevent the firm from benefiting from a decrease in LIBOR. The firm decides that a long position in an FRA is an appropriate response to this risk that it does not wish to bear. There are, however, a few technical details that must be considered.

Recall that the FRA will pay off in 30 days. The loan, however, will be taken out at that time and will be paid back, with interest, 90 days later. Thus, the FRA payoff is designed to protect against interest that is paid 90 days after the FRA expires. To appropriately analyze the effectiveness of the FRA, we might consider the FRA payoff as reducing or increasing the amount borrowed. Alternatively, we could take the FRA payoff, compound it for 90 days, and incorporate it into the amount paid back on the loan. In other words, if the FRA pays a positive amount, we reinvest this amount at LIBOR for 90 days. If it pays a negative amount, we borrow this amount at LIBOR for 90 days. This is the approach we shall take here: we borrow the $20 million in 30 days and compound the FRA payoff 90 days, adding it to or subtracting it from the interest.[3] The results are shown in Table 13.3.

Let us verify one of the outcomes. Let LIBOR be 12 percent. The calculation of the FRA payoff of $97,087 was covered earlier. Compounding this amount for 90 days at 12 percent gives subject to round-off error.

$$\$97{,}087\left(1+0.12\left(\frac{90}{360}\right)\right)=\$100{,}000.$$

[3]It is possible, if not likely, that we might not be able to borrow or invest the payoff at LIBOR, but we shall ignore this point to keep the illustration as simple as possible.

Table 13.3 Hedging an Anticipated Loan with a Forward Rate Agreement

Scenario: A firm plans to borrow $20 million in 30 days at 90-day LIBOR plus 100 basis points. The loan will be paid back with principal and interest 90 days later. Concerned about the possibility of rising interest rates, the firm would like to lock in the rate it pays by going long an FRA. The rate on 30-day FRAs based on 90-day LIBOR is 10%. Interest on the loan and the FRA is based on the factor 90/360. The outcomes for a range of LIBORs are shown below.

LIBOR on Day 30	FRA Payoff on Day 30	FRA Payoff Compounded to Day 120	Amount Due on Loan on Day 120	Total Amount Paid on Day 120	Effective Rate on Loan	Effective Rate without FRA
6.00%	−$197,044	−$200,000	$20,350,000	$20,550,000	11.63%	7.29%
6.50	−172,202	−175,000	20,375,000	20,550,000	11.63	7.82
7.00	−147,420	−150,000	20,400,000	20,550,000	11.63	8.36
7.50	−122,699	−125,000	20,425,000	20,550,000	11.63	8.90
8.00	−98,039	−100,000	20,450,000	20,550,000	11.63	9.44
8.50	−73,439	−75,000	20,475,000	20,550,000	11.63	9.99
9.00	−48,900	−50,000	20,500,000	20,550,000	11.63	10.53
9.50	−24,420	−25,000	20,525,000	20,550,000	11.63	11.08
10.00	0	0	20,550,000	20,550,000	11.63	11.63
10.50	24,361	25,000	20,575,000	20,550,000	11.63	12.18
11.00	48,662	50,000	20,600,000	20,550,000	11.63	12.74
11.50	72,904	75,000	20,625,000	20,550,000	11.63	13.29
12.00	97,087	100,000	20,650,000	20,550,000	11.63	13.85
12.50	121,212	125,000	20,675,000	20,550,000	11.63	14.41
13.00	145,278	150,000	20,700,000	20,550,000	11.63	14.97
13.50	169,287	175,000	20,725,000	20,550,000	11.63	15.54
14.00	193,237	200,000	20,750,000	20,550,000	11.63	16.10

The amount due on the loan is the principal compounded at 12 percent plus the 100 basis point premium for 90 days:

$$\$20,000,000\left(1+(0.12+0.01)\left(\frac{90}{360}\right)\right)=\$20,650,000.$$

The total amount due is $20,650,000 minus the payoff on the FRA of $100,000 for $20,550,000. Thus, the firm borrowed $20 million and 90 days later paid back $20,550,000. The effective rate is found as follows:

$$\left(\frac{\$20,550,000}{\$20,000,000}\right)^{365/90}-1=0.1163$$

Had the firm not done the FRA, it would have paid back $20,650,000. The effective rate without the FRA would then be

$$\left(\frac{\$20,650,000}{\$20,000,000}\right)^{365/90}-1=0.1385$$

We see that the effective rate using the FRA is always 11.63 percent, which is illustrated in Figure 13.2. The higher LIBOR is, the higher the effective rate without the FRA. If LIBOR were below 10 percent, the FRA payoff would be negative and the amount paid back would be larger as a result of the loss on the FRA. But overall, the total amount due would be the same, $20,550,000, and the effective rate would be 11.63 percent.

FRAs have fairly limited uses, being mostly restricted to the situation described here. As noted, they can be used to speculate on interest rates. A series of FRAs with different expirations could be combined to hedge the risk on a floating-rate loan. Since the term structure would not likely be flat, however, each FRA would probably be at a different rate. Thus, the firm would lock in a series of different fixed rates for each of the floating interest payments on its loan. Most firms prefer to lock in the same fixed rate for each of its loan payments. This result is easily achieved with a swap. Indeed a swap is a series of FRAs, but with each

FIGURE 13.2 Cost of Loan with and without FRA

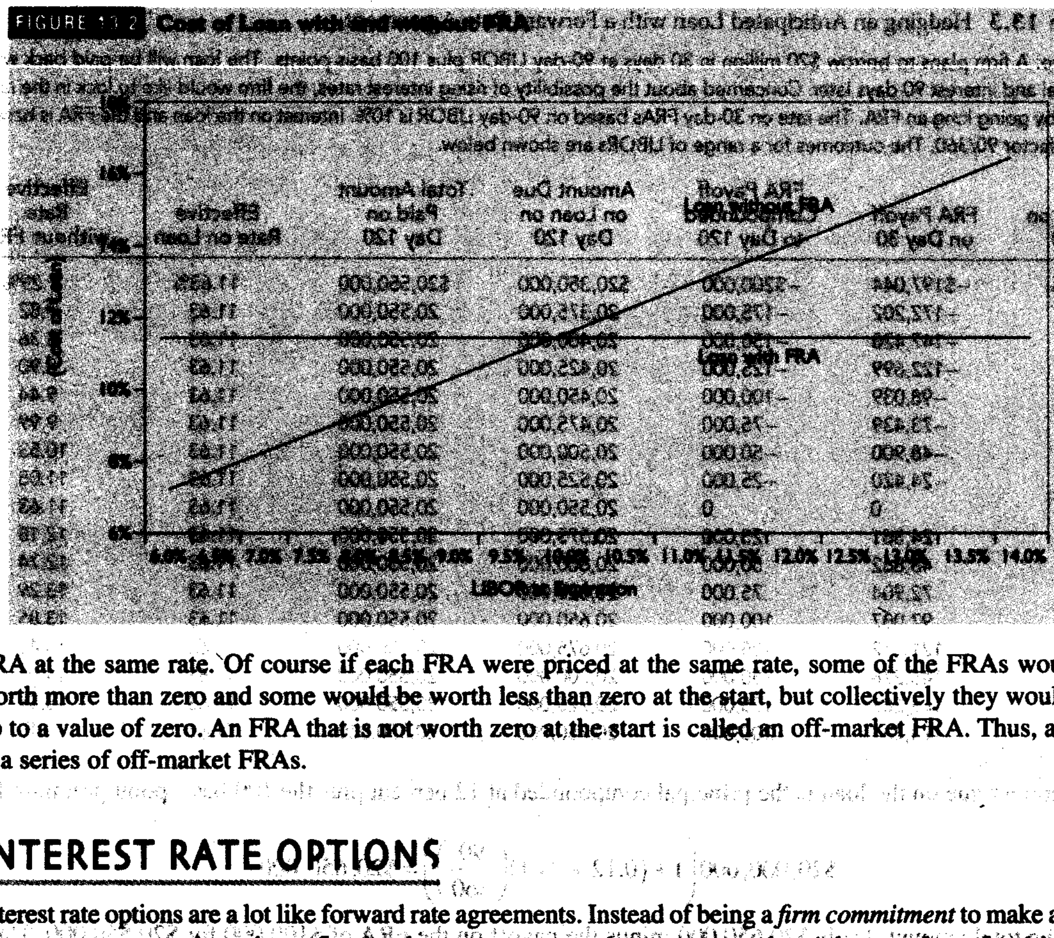

FRA at the same rate. Of course if each FRA were priced at the same rate, some of the FRAs would be worth more than zero and some would be worth less than zero at the start, but collectively they would add up to a value of zero. An FRA that is not worth zero at the start is called an off-market FRA. Thus, a swap is a series of off-market FRAs.

INTEREST RATE OPTIONS

Interest rate options are a lot like forward rate agreements. Instead of being a *firm commitment* to make a fixed interest payment and receive a floating interest payment, they represent the *right* to make a fixed interest payment and receive a floating interest payment or to make a floating interest payment and receive a fixed interest payment. Unlike the options we have already covered, interest rate options have an exercise rate or strike rate, rather than an exercise price or strike price. While European and American versions are available, interest rate options are more often than not of the European variety. This is because they are normally used to hedge an interest rate exposure on a specific date.

Structure and Use of a Typical Interest Rate Option

Like most options, interest rate options come in the form of calls and puts. An interest rate call gives the holder the right to make a known interest payment, based on the exercise rate, and receive an unknown interest payment. This unknown interest payment will usually be determined by an interest rate such as LIBOR. To acquire the option, the buyer pays a premium, the option price, up front. The option expires at a specific date known, of course, as the expiration. As with swaps and FRAs, an interest rate option is based on a given amount of notional principal on which the interest is calculated. Using the symbol X as the exercise rate and m-day LIBOR as the underlying, the payoff of an interest rate call is as follows:

$$(\text{Notional Principal})\,\text{Max}(0, \text{LIBOR} - X)\left(\frac{m}{360}\right)$$

An interest rate put permits the holder the right to pay a floating rate and receive a fixed rate. The payoff of an interest rate put is as follows:

[redacted]

As noted earlier in this chapter, in contrast to the payment on an FRA, the payoff of an interest rate option does not occur at the expiration. If the underlying is m-day LIBOR, the payoff occurs m days after the expiration of the option. Remember that in an FRA, the payoff is made at expiration, but is based on a rate such as m-day LIBOR, which assumes that payment occurs m days later. Hence, the discounting of the payoff is appropriate for an FRA. For an interest rate option (as well as an interest rate swap), the payoff is deferred, so no discounting is required.

Consider interest rate call and put options with notional principals of $20 million, expiring in 30 days; an underlying of 90-day LIBOR; and exercise rates of 10 percent. Let us look at how the payoffs are calculated. Suppose that at expiration LIBOR is 6 percent. The call payoff would be

$$\$20,000,000\left(\text{Max}(0, 0.06 - 0.10)\left(\frac{90}{360}\right)\right) = \$0,$$

and the put payoff would be

$$\$20,000,000\left(\text{Max}(0, 0.10 - 0.06)\left(\frac{90}{360}\right)\right) = \$200,000.$$

If LIBOR at expiration is 14 percent, the call payoff would be

$$\$20,000,000\left(\text{Max}(0, 0.14 - .10)\left(\frac{90}{360}\right)\right) = \$200,000,$$

and the payoff on the put would be

$$\$20,000,000\left(\text{Max}(0, 0.10 - 0.14)\left(\frac{90}{360}\right)\right) = \$0.$$

These payoffs are made 90 days after the expiration of the options.

Pricing and Valuation of Interest Rate Options

In Part I of this book, we devoted a great deal of effort to pricing options. We discussed the binomial model and the Black-Scholes-Merton model for pricing options on assets, and the Black model for pricing options on futures. Pricing interest rate options, however, is much more complicated. To obtain the most accurate interest rate option pricing requires the development of a fairly sophisticated model that will capture movements in the term structure. To capture movements in the term structure, the model must simultaneously reflect movements in all bonds in the market and do so without permitting any arbitrage opportunities. Binomial models of the term structure are commonly used for this purpose, but they are quite sophisticated. We shall approach interest rate option pricing from an alternative and less complex angle. In fact, many interest rate option dealers use this simple approach.

The method we use is the Black model. Remember from Chapter 9 that the Black model is designed for pricing options on futures contracts. Likewise, it can be used for pricing options on forward contracts. The Black model is commonly used in pricing interest rate options by noting that in an interest rate option, the underlying is the current value of the interest rate at expiration. The current value of this rate would be the forward rate. That is, suppose the underlying is 90-day LIBOR and the option expires in 30 days. Then the forward rate for a 90-day Eurodollar time deposit to start in 30 days is the underlying when the option is initiated. That rate evolves over the 30-day life of the option into the 90-day LIBOR at the expiration of the option. Thus, we start by taking the forward rate as the underlying in the Black model. The exercise rate of the option is used as the exercise price in the model. The time to expiration is easily defined as the number of days to expiration of the option divided by 365. The volatility is the volatility of the forward rate. We shall not spend much time on the volatility, relying primarily on our understanding of the general notion of volatility from Chapter 5. Note, however, that we are referring to the volatility of the relative change in a rate, rather than a price. The risk-free rate is the risk-free rate for the period to the option's expiration. It is important to note, however, that the forward rate and risk-free rate should be in continuously compounded form. Technically, the exercise rate should also be viewed as a continuously compounded rate, but we shall just leave it in discrete form. The Black model is not a perfect fit for the situation of pricing an interest rate option, but it provides a reasonable approximation.

The interest rate call option that we shall examine has 30 days until expiration. Thus, T = 30/365 = 0.0822. The exercise rate is 10 percent. The continuously compounded forward rate is obtained as the continuously compounded equivalent of the forward rate based on the term structure. The problem we are working here is the same one we worked when examining FRAs. In Table 13.1 we saw that the 30-day LIBOR is 11 percent, and the 120-day LIBOR is 10.32 percent. Using this information, we calculate the continuously forward rate as follows:

$$\ln\left(\frac{1+0.1032\left(\frac{120}{360}\right)}{1+0.11\left(\frac{30}{360}\right)}\right)\left(\frac{365}{90}\right)=0.1002.$$

The numerator inside the log function is the compound value of $1 invested at 10.32 percent for 120 days. The denominator is the compound value of $1 invested at 11 percent for 30 days. This ratio is one plus the forward rate for 30 days. Taking the log converts it to continuous compounding and multiplying by 365/90 annualizes it.

Finally, we need the continuously compounded risk-free rate for 30 days. We know that $1 invested for 30 days grows to $\$1(1+0.11(30/360)) = 1.0091667$. We take the log of this and multiply by 365/30 to obtain $r = 0.1110$.

Now, suppose we plug these values into the Black model. The result we obtain reflects the assumption that the option payoff occurs at expiration. With interest rate options, the payoff is deferred. Thus, when the underlying is an m-day rate, we need to discount the payoff m days (90 in this problem) at the continuously compounded forward rate, denoted as F. Letting C be the Black call option price:

[redacted]

Finally, we must note that the result we obtain is stated in terms of an annual interest rate. This is an acceptable way to quote the option price, but to obtain the contract premium (the actual amount paid), we must multiply by m/360 and the notional principal:

[redacted]

Table 13.4 illustrates the calculation of the interest rate option price and the contract premium.

Interest Rate Option Strategies

Recall the interest rate call example that we used previously. We purchased a call on 90-day LIBOR expiring in 30 days. The exercise rate was 10 percent. In the previous section, we showed that the price of this call, using a computer for the calculations, would be $17,368. It is easy to consider a scenario in which we might want to buy this call. In fact, this call was created to show an alternative to the FRA as a solution for the problem we discussed earlier in this chapter. Recall that a company was planning to borrow $20 million in

Table 13.4 Calculating an Interest Rate Option Price and Contract Premium Using the Black Model

$F = 0.1002$ $X = 0.10$ $r_c = 0.1110$ $\sigma = 0.3$ $T = 0.0822$ $m = 90$

1. Compute d_1:

$$d_1 = \frac{\ln(0.1002/0.10) + (0.3^2/2)0.0822}{0.3\sqrt{0.0822}} = 0.0662$$

2. Compute d_2:

$$d_1 = 0.0662 - 0.3\sqrt{0.0822} = -0.0198$$

3. Look up $N(d_1)$: $N(0.07) = 0.5279$

4. Look up $N(d_2)$: $N(-0.02) = 1 - N(0.02) = 1 - 0.5080 = 0.4920$

5. Calculate the Black call option price and discount at the risk-free rate for the life of the option to obtain the interest rate option price:

$$C = e^{-0.1110(0.0822)}[0.1002(0.5279) - 0.10(0.4920)] = 0.00366202$$

$$\text{Interest rate option price} = 0.00366202e^{-0.1002(90/365)} = 0.00357265$$

6. Multiply the interest rate option price by the notional principal times m/360:

$$\text{Contract price} = \$20{,}000{,}000\left(\frac{90}{360}\right)0.00357265 = \$17{,}1863$$

Note: using a computer for a more precise result gives a premium of $17,368, which we shall use.

30 days at 90-day LIBOR plus 100 basis points. The loan would involve a single repayment of interest and principal 90 days later. When using an option, however, there is one additional aspect of the problem we must consider. Remember that we pay the option premium today. Thirty days later, the option expires and its payoff is determined. We borrow the money at that point. Ninety days later, we receive the option payout, if any, and pay back the loan. When using FRAs, we had cash flows on day 30 and 90 days later on day 120. When using options we have cash flows on day 0, day 30, and day 120. To determine the effective rate on the loan, we must somehow incorporate the option premium into the cash flows on the loan. We do this by compounding it from day 0 to day 30 at the 30-day rate.[4] This amounts to

$$\$17{,}368\left(1 + 0.11\left(\frac{30}{360}\right)\right) = \$17{,}528.$$

When the loan is taken out, the proceeds of $20 million are effectively only $20,000,000 − $17,528 = $19,982,472, reflecting the fact that the option was purchased in conjunction with the loan. Table 13.5 presents the results of this strategy of hedging a loan with an interest rate call. Let us do a sample calculation.

Suppose LIBOR at expiration is 8 percent. The option payoff is

$$\$20{,}000{,}000\text{Max}(0, 0.08 - 0.10)\left(\frac{90}{360}\right) = \$0.$$

[4] An argument can be made for the fact that the rate used in this calculation should be the 30-day rate plus the 1 percent premium. If the company borrowed this money then it would have to pay the 30-day LIBOR plus the 1 percent premium. Alternatively, if the company used its own cash to buy the option, we should compound the premium at its opportunity cost or the lending rate, which is probably less than LIBOR. For FRAs, this problem is not an issue because no premium is paid up front. There is no clear-cut answer on this question, so we shall just compound it at LIBOR.

The amount due on the loan 90 days later is

$$\$20,000,000\left((1+0.08+0.01)\left(\frac{90}{360}\right)\right)=\$20,450,000.$$

Thus, the total amount due is \$20,450,000. The effective rate on the loan is

$$\left(\frac{\$20,450,000}{\$19,982,472}\right)^{365/90}-1=0.0983.$$

Without the call, the full loan proceeds of \$20 million would be received on day 30. Then the effective rate on the loan would be

$$\left(\frac{\$20,450,000}{\$20,000,000}\right)^{365/90}-1=0.0944.$$

Table 13.5 Hedging an Anticipated Loan with an Interest Rate Call

Scenario: A firm plans to borrow \$20 million in 30 days at 90-day LIBOR plus 100 basis points. The loan will be paid back with principal and interest 90 days later. Concerned about the possibility of an increase in interest rates before the loan is taken out, the firm would like to protect against an increase in interest rates, while preserving the flexibility to benefit from a decrease in interest rates. An interest rate call is appropriate for this objective. Interest on the loan and the call is based on the factor 90/360. The call will have an exercise rate of 10 percent and cost \$17,368. The cost of the call is compounded for 30 days at the 30-day rate of 11 percent to obtain \$17,528, which is the effective cost of the call at the time the loan is taken out. The proceeds from the loan are \$20,000,000 – \$17,528 = \$19,982,472. The outcomes for a range of LIBORs are shown below.

LIBOR on Day 30	Call Payoff on Day 120	Amount Due on Loan on Day 120	Total Amount Paid on Day 120	Effective Rate on Loan	Effective Rate without Call
6.0%	\$0	\$20,350,000	\$20,350,000	7.67%	7.29%
6.5	0	20,375,000	20,375,000	8.21	7.82
7.0	0	20,400,000	20,400,000	8.75	8.36
7.5	0	20,425,000	20,425,000	9.29	8.90
8.0	0	20,450,000	20,450,000	9.83	9.44
8.5	0	20,475,000	20,475,000	10.38	9.99
9.0	0	20,500,000	20,500,000	10.92	10.53
9.5	0	20,525,000	20,525,000	11.48	11.08
10.0	0	20,550,000	20,550,000	12.03	11.63
10.5	25,000	20,575,000	20,550,000	12.03	12.18
11.0	50,000	20,600,000	20,550,000	12.03	12.74
11.5	75,000	20,625,000	20,550,000	12.03	13.29
12.0	100,000	20,650,000	20,550,000	12.03	13.85
12.5	125,000	20,675,000	20,550,000	12.03	14.41
13.0	150,000	20,700,000	20,550,000	12.03	14.97
13.5	175,000	20,725,000	20,550,000	12.03	15.54
14.0	200,000	20,750,000	20,550,000	12.03	16.10

Now, suppose LIBOR at expiration is 12 percent. The option payoff is

$$\$20,000,000\text{Max}(0,0.12-0.10)\left(\frac{90}{360}\right)=\$100,000.$$

The amount due on the loan 90 days later is

$$\$20,000,000\left(1+(0.12+0.01)\left(\frac{90}{360}\right)\right)=\$20,650,000.$$

Thus, the total amount due is \$20,650,000 less the option payoff of \$100,000 for a total of \$20,550,000. The effective rate on the loan is

$$\left(\frac{\$20{,}550{,}000}{\$19{,}982{,}472}\right)^{365/90} - 1 = 0.1203.$$

Without the call, the full loan proceeds of \$20 million would be received on day 30. Then the effective rate on the loan would be

$$\left(\frac{\$20{,}650{,}000}{\$20{,}000{,}000}\right)^{365/90} - 1 = 0.1385.$$

FIGURE 13.3 Cost of Loan with and without Interest Rate Call

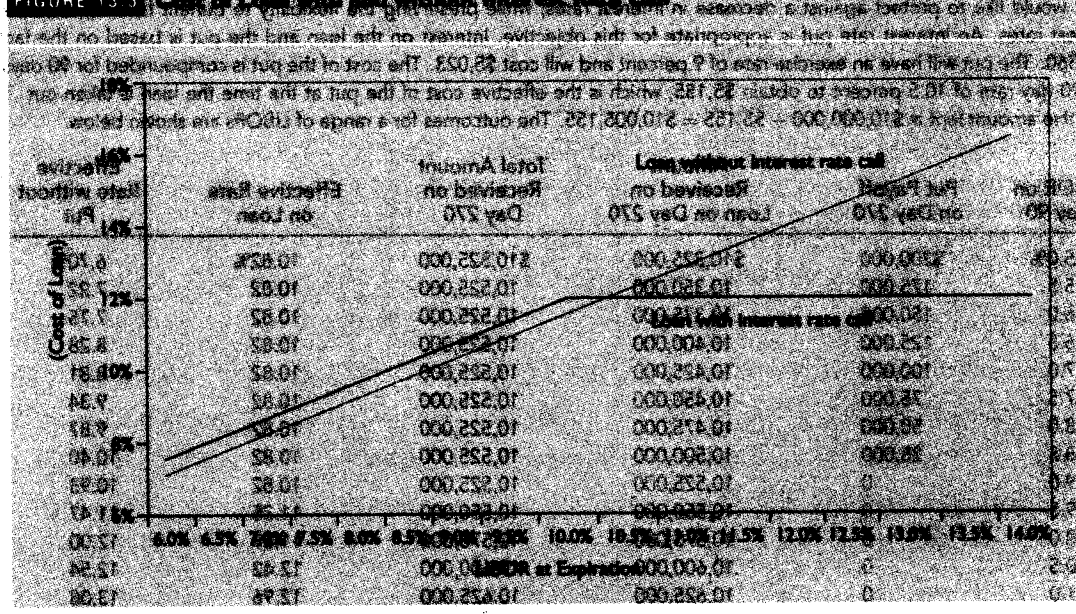

Figure 13.3 illustrates the effective rate on the loan with and without the call. Observe how the effective rate increases without limit in the absence of the call. The call limits the effective rate. If the 30-day LIBOR at expiration is below the exercise rate of 10 percent, however, the call is not exercised and the cost of the call raises the effective rate on the loan.

Let us look at an interest rate put in a different scenario. Consider a bank that plans to make a \$10 million floating-rate loan in 90 days. The loan will be for 180 days, and the rate will be 180-day LIBOR plus 150 basis points. The current 90-day LIBOR is 10.5 percent. The bank is worried about falling interest rates over the period from now until the loan starts. Thus, an interest rate put option would be appropriate. An interest rate put option with an exercise rate of 9 percent is priced at \$5,023 for this contract.[5] This premium would then be compounded for 90 days at the 90-day rate of 10.5 percent to obtain a value of

$$\$5{,}023\left(1 + 0.105\left(\frac{90}{360}\right)\right) = \$5{,}155.$$

[5]You can verify this result using the Black model for an interest rate put and the following information. The continuously compounded forward rate is 9.45 percent; the exercise rate is 9 percent, the continuously compounded risk-free rate is 10.51 percent, the time to expiration is 90/365 = 0.2466, and the volatility is 0.15. The answer of \$5,023 was obtained on a computer. Your answer computed by hand should be about \$4,323.

Even though the option is purchased on day 0, this is the effective cost of the option on the day the loan is taken out, 90 days later. Thus, when the bank makes the loan it effectively pays out $10,000,000 + $5,155 = $10,005,155. The results are presented in Table 13.6.

Suppose LIBOR at expiration is 5 percent. Then the payoff of the put is

$$\$10{,}000{,}000\text{Max}(0, 0.09 - 0.05)\left(\frac{180}{360}\right) = \$200{,}000.$$

Table 13.6 Hedging an Anticipated Loan with an Interest Rate Put

Scenario: A bank plans to lend $10 million in 90 days at 180-day LIBOR plus 150 basis points. The loan will be paid back with principal and interest 180 days later. Concerned about the possibility of falling interest rates before the loan is taken out, the bank would like to protect against a decrease in interest rates, while preserving the flexibility to benefit from an increase in interest rates. An interest rate put is appropriate for this objective. Interest on the loan and the put is based on the factor 180/360. The put will have an exercise rate of 9 percent and will cost $5,023. The cost of the put is compounded for 90 days at the 90-day rate of 10.5 percent to obtain $5,155, which is the effective cost of the put at the time the loan is taken out. The effective amount lent is $10,000,000 + $5,155 = $10,005,155. The outcomes for a range of LIBORs are shown below.

LIBOR on Day 90	Put Payoff on Day 270	Amount Received on Loan on Day 270	Total Amount Received on Day 270	Effective Rate on Loan	Effective Rate without Put
5.0%	$200,000	$10,325,000	$10,525,000	10.82%	6.70%
5.5	175,000	10,350,000	10,525,000	10.82	7.22
6.0	150,000	10,375,000	10,525,000	10.82	7.75
6.5	125,000	10,400,000	10,525,000	10.82	8.28
7.0	100,000	10,425,000	10,525,000	10.82	8.81
7.5	75,000	10,450,000	10,525,000	10.82	9.34
8.0	50,000	10,475,000	10,525,000	10.82	9.87
8.5	25,000	10,500,000	10,525,000	10.82	10.40
9.0	0	10,525,000	10,525,000	10.82	10.93
9.5	0	10,550,000	10,550,000	11.35	11.47
10.0	0	10,575,000	10,575,000	11.89	12.00
10.5	0	10,600,000	10,600,000	12.42	12.54
11.0	0	10,625,000	10,625,000	12.96	13.08
11.5	0	10,650,000	10,650,000	13.50	13.62
12.0	0	10,675,000	10,675,000	14.04	14.16
12.5	0	10,700,000	10,700,000	14.59	14.71
13.0	0	10,725,000	10,725,000	15.13	15.25

The amount due on the loan is

$$\$10{,}000{,}000\left(1 + (0.05 - 0.015)\left(\frac{180}{360}\right)\right) = \$10{,}325{,}000.$$

So the bank receives $10,325,000 on the loan plus $200,000 on the put for a total amount due of $10,525,000. As we noted, the effective outlay on the loan is $10,005,155. Thus, the effective rate on the loan is

$$\left(\frac{\$10{,}525{,}000}{\$10{,}005{,}155}\right)^{365/180} - 1 = 0.1082$$

Without the put, the outlay on the loan is only $10,000,000, but the effective rate on the loan is

$$\left(\frac{\$10{,}325{,}000}{\$10{,}000{,}000}\right)^{365/180} - 1 = 0.067.$$

Suppose LIBOR at expiration is 13 percent. Then the payoff of the put is

$$\$10,000,000\text{Max}(0,0.09-0.13)\left(\frac{180}{360}\right)=\$0.$$

The amount due on the loan is

$$\$10,000,000\left(1+(0.13+0.015)\left(\frac{180}{360}\right)\right)=\$10,725,000.$$

So the bank receives \$10,725,000 on the loan but nothing on the put. Given the effective outlay on the loan of \$10,005,155, the effective rate on the loan is

$$\left(\frac{\$10,725,000}{\$10,005,155}\right)^{365/180}-1=0.1513.$$

Without the put, the effective rate on the loan is

$$\left(\frac{\$10,725,000}{\$10,000,000}\right)^{365/180}-1=0.1525.$$

A graph of the results is in Figure 13.4. Note how the loan without the put has considerable downside risk. The loan with the put participates in the benefits from high interest rates but is protected against low interest rates. If interest rates are higher, however, the loan generates a lower return, due to the put premium, than the loan without the put.

Interest Rate Caps, Floors, and Collars

What we have covered so far with respect to interest rate options does not include the majority of the situations in which interest rate options are used. Recall from Chapter 12 how we showed that swaps are widely used in conjunction with floating-rate loans. In this chapter so far, we have examined only loans involving one payment. FRAs, interest rate calls, and interest rate puts are not useful for loans involving more than one payment, at least not in the form we have seen them here. Interest rate calls and puts can, however,

FIGURE 13.4 Return on Loan with and without Interest Rate Put

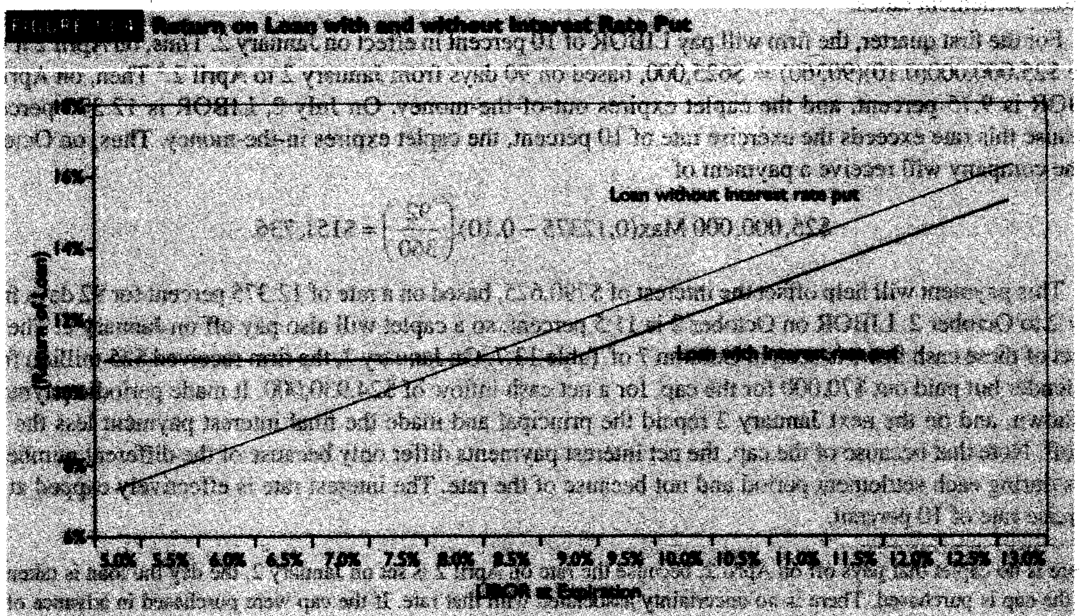

be combined into a series of options that can effectively protect floating-rate loans. The analogy is that a swap, which is effectively a combination of FRAs, can protect a floating-rate loan, as we saw in Chapter 12. Of course, swaps and FRAs are commitments. In some cases, an option is preferred. Options provide flexibility to benefit from rate movements in one direction, while not being hurt by rate movements in another direction.

A combination of interest rate calls designed to protect a borrower in a floating-rate loan against increases in interest rates is called an interest rate cap. Each component call is referred to as a caplet. A combination of interest rate puts designed to protect a lender in floating-rate loan against decreases in interest rates is called an interest rate floor. Each component put is referred to as a floorlet. A combination of a long cap and short floor is called an interest rate collar, which is similar to but still different from the collars we discussed in Chapter 7. A collar is most often used by a borrower and consists of a long position in a cap, financed by selling a short position in a floor.

Interest Rate Caps Let us first look at an example of a cap. On January 2, a firm borrows $25 million over one year. It will make payments on April 2, July 2, October 2, and next January 2. On each of these dates, starting with January 2, LIBOR in effect on that date will be the interest rate paid over the next three months. The current LIBOR is 10 percent. The firm wishes to fix the rate on each payment at no more than 10 percent, so it buys a cap for an up-front payment of $70,000 with an exercise rate of 10 percent. The payoffs are based on the exact number of days in the settlement period and a 360-day year. At each interest payment date, the cap will be worth

$$\$25{,}000{,}000\,\text{Max}(0, \text{LIBOR} - 0.10)\left(\frac{\text{days}}{360}\right),$$

where LIBOR is understood to be set on the previous settlement date and "days" is the number of days in the period. LIBOR for the first payment is set when the loan is initiated. Thus, as we previously discussed for interest rate options, the decision to exercise is made at the beginning of the settlement period, which is the date on which the rate is set, but the payoff occurs at the end of the settlement period. If LIBOR exceeds 10 percent, the firm will exercise the option and receive an amount as given by this equation. This payoff helps offset the higher interest rate on the loan. Table 13.7 illustrates a set of possible payments that might occur on this loan. Of course, these payments represent only one of an infinite number of possible LIBORs on the various settlement dates.

For the first quarter, the firm will pay LIBOR of 10 percent in effect on January 2. Thus, on April 2 it will owe $25,000,000(0.10)(90360) = $625,000, based on 90 days from January 2 to April 2.[6] Then, on April 2, LIBOR is 9.75 percent, and the caplet expires out-of-the-money. On July 2, LIBOR is 12.375 percent. Because this rate exceeds the exercise rate of 10 percent, the caplet expires in-the-money. Thus, on October 2, the company will receive a payment of

$$\$25{,}000{,}000\,\text{Max}(0{,}12375 - 0.10)\left(\frac{92}{360}\right) = \$151{,}736.$$

This payment will help offset the interest of $790,625, based on a rate of 12.375 percent for 92 days from July 2 to October 2. LIBOR on October 2 is 11.5 percent, so a caplet will also pay off on January 2. The net effect of these cash flows is seen in Column 7 of Table 13.7. On January 2, the firm received $25 million from the lender but paid out $70,000 for the cap, for a net cash inflow of $24,930,000. It made periodic payments as shown, and on the next January 2 repaid the principal and made the final interest payment less the cap payoff. Note that because of the cap, the net interest payments differ only because of the different number of days during each settlement period and not because of the rate. The interest rate is effectively capped at the exercise rate of 10 percent.

[6]There is no caplet that pays off on April 2, because the rate on April 2 is set on January 2, the day the loan is taken out and the cap is purchased. There is no uncertainty associated with that rate. If the cap were purchased in advance of the day on which the loan is taken out, the firm might add a caplet that would expire on January 2 and pay off on April 2.

TABLE 13.7 After-the-Fact Payments for Loan with Interest Rate Cap

Scenario: On January 2, a company takes out a $25 million one-year loan with interest paid quarterly at LIBOR. To protect against rising interest rates, the company buys an interest rate cap with an exercise rate of 10 percent for a premium of $70,000. The interest payments on the loan and the payoffs of the cap are based on the exact number of days and a 360-day year.

Date	Days in Period	LIBOR (%)	Interest Due	Cap Payment	Principal Repayment	Net Cash Flow	Net Cash Flow without Cap
Jan 2		10.000		−$70,000	$0	$24,930,000	$25,000,000
Apr 2	90	9.750	$625,000	—	0	−625,000	−625,000
Jul 2	91	12.375	616,146	0	0	−616,146	−616,146
Oct 2	92	11.500	790,625	151,736	0	−638,889	−790,625
Jan 2	92		734,722	95,833	25,000,000	−25,638,889	−25,734,722

Effective annual rate
Without cap: 11.50%
With cap: 10.78%

Note: This combination of LIBORs on the above dates represents only one of an infinite number of possible outcomes. They are used only to illustrate how the payments are determined and not the likely results.

If we wish to know the annualized rate the firm actually paid, we must solve for the internal rate of return, which requires a computer or financial calculator. We are solving for the rate that equates the present value of the four payouts to the initial receipt:

$$\$24,930,000 = \frac{\$625,000}{(1+y)^1} + \frac{\$616,146}{(1+y)^2} + \frac{\$638,889}{(1+y)^3} + \frac{\$25,638,889}{(1+y)^4}.$$

The solution is $y = 0.025927$. Annualizing this result gives a rate of $(1.025927)^4 - 1 = 0.1078$. The last column in Table 13.7 shows the cash flows if the cap had not been purchased. Solving for the internal rate of return using those numbers and annualizing this result gives a rate of 0.115. Thus, the cap saved the firm 72 basis points, because during the life of the loan, interest rates were higher than they were at the time the loan was initiated. Of course, if rates had fallen, the firm probably would have paid a higher effective rate with the cap, because the caplets would have been out-of-the-money but the premium was expended.

Pricing caps proceeds in the same manner as pricing interest rate calls. Each caplet is a separate interest rate call, but with a different expiration and a potentially different forward rate, risk-free rate, and volatility. The total price of the cap is the sum of the prices of the component caplets.

Interest Rate Floors An interest rate floor is typically used by a lender in a floating-rate loan who wants protection against falling rates. As noted above, a floor contains a series of interest rate put options, each of which is called a floorlet. At each interest payment date, the payoff of an interest rate floor tied to LIBOR with an exercise rate of, say, 8 percent, payoffs based on the exact number of days and a 360-day year, and a notional principal of $15 million will be

$$\$15,000,000\text{Max}(0, 0.08 - \text{LIBOR})\left(\frac{\text{days}}{360}\right),$$

where as previously noted, LIBOR is set at the beginning of the settlement period.

Suppose on December 16 a bank makes a one-year, $15 million loan with payments at LIBOR on March 16, June 16, September 15, and next December 16. LIBOR is currently 7.875 percent. Thus, on March 16 the bank will receive $15,000,000 [0.07875(90/360)] = $295,313 in interest, which is based on 90 days between December 16 and March 16. The bank purchases a floor for $30,000 at an exercise rate of 8 percent. The new rate on March 16 is 8.25 percent, so the first floorlet expires out-of-the-money. LIBOR on June 16 is 7.125 percent, however, so the floorlet expires in-the-money and pays off

$$\$15{,}000{,}000\text{Max}(0, 0.08 - 0.07125)\left(\frac{91}{360}\right) = \$33{,}177,$$

on the next interest payment date of September 15. This payoff will add to the interest received of $270,156, which is lower because of the fall in interest rates. The complete results for this one-year loan under a series of assumed LIBORs over the life of the loan are shown in Table 13.8.

The bank paid out $15,000,000 up front to the borrower and another $30,000 for the floor. Column 7 in Table 13.8 indicates the periodic cash flows associated with the loan combined with the floor. Following the same procedure as in the cap, we can solve for the periodic rate that equates the present value of the inflows to the outflow. This rate turns out to be about 1.9831 percent. Annualizing gives a rate of $(1.019831)^4 - 1 = .0817$. The last column shows the cash flows if the floor had not been used. The annualized return without the floor is 7.41 percent. Thus, the floor boosted the bank's return by 76 basis points. Of course, in a period of rising rates, the bank will gain less from the increase in interest rates because the floor premium is lost but the floorlets will tend to not be exercised, in which case the bank might have been better off without the floor.

TABLE 13.8 After-the-Fact Payments for Loan with Interest Rate Floor

Scenario: On December 16, a bank makes a $15 million one-year loan with interest paid quarterly at LIBOR. To protect against falling interest rates, the bank buys an interest rate floor with an exercise rate of 8 percent for a premium of $30,000. The interest payments on the loan and the payoffs of the floor are based on the exact number of days and a 360-day year.

Date	Days in Period	LIBOR (%)	Interest Due	Floor Payment	Principal Repayment	Net Cash Flow	Net Cash Flow without Floor
Dec 16		7.875		−$30,000	$0	−$15,030,000	−$15,030,000
Mar 16	90	8.250	$295,313	—	0	295,313	295,313
Jun 16	92	7.125	316,250	0	0	316,250	316,250
Sep 15	91	6.000	270,156	33,177	0	303,333	270,156
Dec 16	92		230,000	76,667	$15,000,000	15,306,667	15,230,000

Effective annual rate
Without floor: 7.41%
With floor: 8.17%

Note: This combination of LIBORs on the above dates represents only one of an infinite number of possible outcomes. They are used only to illustrate how the payments are determined and not the likely results.

As with caps, the price of a floor is found by pricing each of the component floorlets and adding up their prices.

Interest Rate Collars In Chapter 7 we discussed how an investor could use a collar to protect a portfolio of stock. The investor owns the portfolio and purchases a put for downside protection. To finance the purchase of the put, the investor sells a call in which the exercise price of the call is set at a level that will produce a premium level on the call that will offset the premium paid for the put. This transaction creates a range for the value of the stock. The value of the position will not fall below the exercise price of the long put, nor will it rise above the exercise price of the short call. In a similar manner, a borrower can use a collar to create a range of interest rates, ensuring that the borrower will pay no more than a given rate, nor any less than another rate.

Consider a firm planning to borrow money that decides to purchase an interest rate cap. In doing so, the firm is trying to place a ceiling (cap) on the rate it will pay on the loan. If rates fall, it can gain by paying a lower rate. In some cases, however, a firm will find it more advantageous to give up the right to gain from falling rates in order to lower the cost of the cap. One way to do this is to sell a floor. The combination of a long interest rate cap and short interest rate floor is called an *interest rate collar,* or sometimes just a collar. While it is not necessary that the premium from selling the floor be exactly equal to the premium from buying the cap, in most situations that is what borrowers prefer. This type of collar is called a zero-cost collar. This term is somewhat misleading, however, because nothing comes at zero cost. The cost to the borrower is in its willingness to give up the benefits from a decrease in interest rates below the exercise price of the floor. A collar establishes a range of rates within which there is interest rate uncertainty, but the maximum and minimum rates are locked in.

Table 13.9 illustrates a zero cost collar in which a firm borrowing $50 million over two years buys a cap for $250,000 with an exercise rate of 10 percent and sells a floor for $250,000 with an exercise rate of 8.5 percent. The loan begins on March 15 and will require payments at approximately 91-day intervals at LIBOR.

By now you should be able to verify the numbers in the table. The interest paid on June 15 is based on LIBOR on March 15 of 10.5 and 92 days during the period. The cap pays off on September 14 and December 14 because those are the ends of the settlement periods in which LIBOR at the beginning of the period was greater than 10 percent. The floor pays off on September 14 and December 15 of the next year and on March 14 of the following year, the due date on the loan. Note that when the floor pays off, the firm makes rather than receives the payment.

Column 8 of Table 13.9 shows the cash flows associated with the collar. Following the procedure previously described to solve for the internal rate of return gives an annualized rate of 9.83 percent for the loan with the collar. Column 9 shows the cash flows had the firm done only the loan with the cap. The rate associated with this strategy would have been 9.93 percent. The last column indicates the cash flows had the firm done neither the cap nor the floor. The borrowing rate associated with that strategy would have been 10.10 percent.

The cap by itself would have helped lower the firm's cost of borrowing. By selling the floor, and thus creating a collar, the cost of the loan was lowered from 9.93 percent to 9.83 percent. Other outcomes, of course, would lead to different results, and it is possible that the loan without the cap or collar would end up best.

TABLE 13.9 After-the-Fact Payments for Loan with Interest Rate Collar

Scenario: On March 15, a company takes out a $50 million two-year loan with interest paid quarterly at LIBOR. To protect against rising interest rates, the company buys an interest rate cap with an exercise rate of 10 percent for a premium of $250,000. To offset the cost of the cap, the company sells an interest rate floor with an exercise rate of 8.5 percent for $250,000. The interest payments on the loan and the payoffs of the cap and floor are based on the exact number of days and a 360-day year.

Date	Days in Period	LIBOR (%)	Interest Due	Cap Payment	Floor Payment	Principal Repayment	Net Cash Flow with Collar	Net Cash Flow with Cap Only	Net Cash Flow without Cap or Floor
Mar 15		10.500		−$250,000	$250,000	$0	$50,000,000	$49,750,000	$50,000,000
Jun 15	92	11.500	$1,341,667	—	—	0	−1,341,667	−1,341,667	−1,341,667
Sep 14	91	11.750	1,453,472	189,583	0	0	−1,263,889	−1,263,889	−1,453,472
Dec 14	91	9.125	1,485,069	221,181	0	0	−1,263,889	−1,263,889	−1,485,069
Mar 15	91	9.500	1,153,299	0	0	0	−1,153,299	−1,153,299	−1,153,299
Jun 14	91	7.625	1,200,694	0	0	0	−1,200,694	−1,200,694	−1,200,694
Sep 14	92	8.375	974,306	0	−111,806	0	−1,086,111	−974,306	−974,306
Dec 15	92	8.000	1,070,139	0	−15,972	0	−1,086,111	−1,070,139	−1,070,139
Mar 14	89		988,889	0	−61,806	50,000,000	−51,050,694	−50,988,889	−50,988,889

Effective annual rate
Without collar: 10.10%
With collar: 9.83%
With cap only: 9.93%

Note: This combination of LIBORs on the above dates represents only one of an infinite number of possible outcomes. They are used only to illustrate how the payments are determined and not the likely results.

Collars are most often used by borrowers, but they could certainly be used by lenders. A lender could buy a floor to protect against falling rates below the floor exercise rate and sell a cap that would give up the gains from rising rates above the cap exercise rate.

Interest Rate Options, FRAs, and Swaps

Before we leave interest rate options, it is important for us to take a look at the relationships between the three most important interest rate derivatives: swaps, FRAs, and options. Recall that in this chapter we first introduced the FRA, which is a forward contract on an interest rate. A party agrees to pay a fixed rate and receive a floating rate. We showed that the fixed rate is the forward rate in the term structure. A swap is like a series of FRAs, but where all fixed payments are at the same rate. In other words, it is a series of FRAs, but each FRA is priced at the swap fixed rate. That means that some of the component FRAs will be priced at a fixed rate higher than they would have if priced individually. Some will have a lower fixed rate. As we mentioned earlier, when an FRA is not priced at the forward rate, it is called an off-market FRA. Thus, a swap is a series of off-market FRAs.

A combination of interest rate calls and puts can also be shown to reproduce the payments on a swap. Assume that we buy a call (in this context, also known as a caplet) and sell a put (floorlet). Let R be the fixed rate on a swap and let X be the exercise rate on both a call and a put. Let the underlying rate be LIBOR. Suppose that we are at a settlement date. The long call payoff is

0	if	$\text{LIBOR} \leq X$
$\text{LIBOR} - X$	if	$\text{LIBOR} > X$

and the short put payoff is

$-(X - \text{LIBOR})$	if	$\text{LIBOR} \leq X$
0	if	$\text{LIBOR} > X$.

Thus, the combined payoff of the long call and short put is

$\text{LIBOR} - X$	if	$\text{LIBOR} \leq X$
$\text{LIBOR} - X$	if	$\text{LIBOR} > X$.

This result can be simplified to

$$\text{LIBOR} - X,$$

which will occur on the next settlement date. This result applies to each settlement date.

Now let us examine the payoffs of a swap to pay a fixed rate R and receive a floating rate. The payoff is

$$\text{LIBOR} - R,$$

which will occur on the next settlement date. This result applies to each settlement date.

The swap and long call/short put results are similar. Whether they are the same depends on whether X = R. Remember from Chapter 12 that we found R by pricing the swap, using the information in the term structure. The exercise rate on a call or put is chosen by the parties to the contract. Thus, X is arbitrary and can be any rate. When X is chosen to equal R, the long call/short put produces the same payoffs as the pay-fixed, receive-floating swap. Recall that a swap requires no initial outlay. Thus, the value of the swap is zero at the start. Hence, when X is chosen to equal R, the long call/short put must also have a zero value at the start. This means that the premium on the call must equal the premium on the put. Otherwise, there would be an arbitrage opportunity.

Since a swap has multiple payments, there would have to be a series of calls and puts. It is important to note, however, that the option strategy we have just described is not a zero-cost collar. We noted that in a zero-cost collar, the call and put premiums are the same but the exercise prices are different. In fact, it is not a collar at all because the options have the same exercise prices.

DERIVATIVES TOOLS

Concepts, Applications, and Extensions

Binomial Pricing of Interest Rate Options

We noted in this chapter that the Black model provides a simple approximation for interest rate option prices. A better approach is to build a model of the term structure that eliminates the possibility of earning arbitrage profits by trading bonds or derivatives with prices driven by interest rates. Binomial models are commonly used to price interest rate options. Here we look at a simple binomial model of the term structure and use it to price an interest rate option.

Consider the term structure of LIBORs below.

Term	Rate	Discount Bond Price
360 days	8.0%	$B_0(360) = 1/(1 + 0.08(360/360)) = 0.9259$
720 days	8.5%	$B_0(720) = 1/(1 + 0.085(720/360)) = 0.8547$

Let us define one-period binomial period as 360 days. The information above will be sufficient to build a two-period binomial model. At time zero, we have the prices of a one-period bond (0.9259) and a two-period bond (0.8547). At time 1 in either outcome, the original one-period will mature to its face value of 1. At time 1 the original two-period bond will be worth a price denoted as $B_1(2)^+$ in the up state or $B_1(2)^-$ in the down state. Thus, the binomial term structure will look like this:

$$B_1(1)^+ = 1$$
$$B_1(2)^+ = ?$$

$$B_0(1) = 0.9259$$
$$B_0(2) = 0.8547$$

$$B_1(1)^- = 1$$
$$B_1(2)^- = ?$$

We need to find the missing bond prices and convert them to interest rates. It is very important to note that the bond prices in the higher states are bond prices in higher interest rate states. Thus, $B_1(2)^+$ is less than $B_1(2)^-$, which may seem a little counterintuitive. That is, the plus (1) state at time 1 is a higher interest rate but lower bond price than the minus (–) state.

To eliminate arbitrage opportunities, we must impose the following restrictions:

- It is not possible to make an arbitrage profit by buying a one-period bond, financing it by selling short a two-period bond.
- It is not possible to make an arbitrage profit by buying a two-period bond, financing it by selling short a one-period bond.

For the first strategy, an arbitrage profit would occur if

$$\frac{1}{B_0(1)} > \frac{B_1(2)^-}{B_0(2)}.$$

An arbitrage profit would occur if the return from investing \$1 in a one-period bond, the left-hand side above, is higher than the highest return one would owe from shorting a two-period bond, the right-hand side above. For the second strategy, an arbitrage profit would occur if

$$\frac{1}{B_0(1)} < \frac{B_1(2)^+}{B_0(2)}.$$

An arbitrage profit would occur if the financing cost from shorting a one-period bond, the left-hand side above, were less than the worst outcome from buying a two-period bond, the right-hand side above. To prevent both conditions from holding, we require that

$$\frac{B_1(2)^+}{B_0(2)} < \frac{1}{B_0(1)} < \frac{B_1(2)^-}{B_0(2)}.$$

Although there are many ways to meet this condition, a simple one is to apply weights of 0.5 and 0.5 to the two expressions on the ends of the inequality, giving us

$$0.5B_1(2)^+ + 0.5B_1(2) = \frac{B_1(2)}{B_0(1)}.$$

The right-hand side is the forward price from the term structure of one- and two-period bonds. The left-hand side can be interpreted as the expected price of the two-period bond at time 1 for a probabilities of 0.5.

This requirement is a single equation with two unknowns, $B_1(2)^+$ and $B_1(2)^-$, so we cannot solve it. We can, however, solve it by adding the constraint that the volatility of the interest rate is a particular known value, σ. Now, we have two equations and two unknowns. Let r^+ and r^- be the two possible continuously compounded one-period rates at time 1. The volatility of the interest rate in a binomial model is found as follows:

$$\sigma = \frac{r^+ + r^-}{2}.$$

These two rates determine the prices, $B_1(2)^+$ and $B_1(2)^-$. Although solving these two equations for the bond prices is not simple, it can be done. The results are

$$r^+ = 0.0901 \qquad r^- = 0.0701,$$

which correspond to zero coupon bond prices of

$$B_1(2)^+ = e^{-0.0901} = 0.9138 \; B_1(2)^- = e^{-0.0701} = 0.9323.$$

The LIBORs that produce these prices are 9.43 percent and 7.26 percent as verified in the following:

$$B_1(2)^+ = \frac{1}{1 + 0.0943\left(\frac{360}{360}\right)} = 0.9138.$$

$$B_1(2)^- = \frac{1}{1 + 0.0726\left(\frac{360}{360}\right)} = 0.9323.$$

Thus, the one-period rates at time 1 are 9.43 percent and 7.26 percent.

As an example of how this model would be used, the payoff of a call option payoff expiring at time 1 with an exercise rate of 8 percent would be

$$C^+ = \text{Max}(0, 0.0943 - 0.08)0.9138 = 0.0131$$
$$C^- = \text{Max}(0, 0.0726 - 0.08)0.9323 = 0.0,$$

where we multiply by the one-period bond price to discount for the delay in the payoff. The option price at time 0 based on a \$1 notional principal is, therefore,

$$C = (0.5(0.0131) + 0.5(0.0))0.9323 = 0.0061,$$

where we see that the weighted average of the next two prices is discounted using the one-period discount factor.

To price options with longer expirations we would need information on the prices of longer-term zero coupon bonds and the volatilities of other interest rates.

INTEREST RATE SWAPTIONS AND FORWARD SWAPS

Earlier in this chapter, we mentioned that a firm that had entered into a swap could exit the swap before it terminates if the firm had previously purchased an option on the swap. Options on swaps are called swap options or *swaptions*. A swaption is an option in which the buyer pays a premium up front and acquires the right to enter into a swap. Although swaptions exist on interest rate, currency, equity, and commodity swaps, we shall focus exclusively on swaptions on interest rate swaps, where it is understood that the underlying is a plain vanilla swap. In an interest rate swaption, the buyer receives the right to enter into a swap as either a fixed-rate payer, floating-rate receiver or as a floating-rate payer, fixed-rate receiver. The former is called a payer swaption, while the latter is called a receiver swaption.

A swaption can be viewed as a variation of an interest rate option, but there are several important differences. Like an interest rate option, the exercise rate is stated in terms of an interest rate. A swaption is based on an underlying swap and has a fixed maturity. A swaption can be American- or European-style. The parties agree up front that exercise will be accomplished by entering into the underlying swap or by an equivalent cash settlement.

Structure of a Typical Interest Rate Swaption

Consider the following situation. A company called MPK Resources is considering the possibility that it will need to engage in an interest rate swap two years from now with a notional principal of $10 million. It expects that the swap would be a three-year pay fixed-receive floating swap. The firm is concerned about rising interest rates over the next two years that would force it to pay a higher fixed rate if it entered into the swap at that time. It thus decides to purchase a two-year European-style payer swaption where the underlying is a three-year, pay fixed-receive floating swap. Naturally the underlying swap should be identical to the one MPK expects to take out in two years. MPK specifies an exercise rate of 11.5 percent. MPK pays a premium up front, the amount of which we do not need to know at this point. To keep the illustration as simple as possible, assume that the underlying swap calls for annual interest payments.

Now let us consider what happens when the swaption expires in two years. At that time we observe a term structure of LIBOR as follows:

Term	Rate	Discount Bond Price
360 days	12.00%	$B_0(360) = 1/(1 + 0.12(360/360)) = 0.8929$
720 days	13.28%	$B_0(720) = 1/(1 + 0.1328(720/360)) = 0.7901$
1080 days	14.51%	$B_0(1080) = 1/(1 + 0.1451(1080/360)) = 0.6967$

Based on this information, what would be the rate in the market for three-year swaps? This result is easily found using what we learned in Chapter 12:

$$R = \left(\frac{1 - 0.6967}{0.8929 + 0.7901 + 0.6967}\right)\left(\frac{360}{360}\right) = 0.1275.$$

Thus, MPK, or any party, could enter into a pay-fixed, receive-floating swap at a fixed rate of 12.75 percent. As it turns out, however, MPK holds a payer swaption that allows it to enter into a pay-fixed, receive-floating swap paying a fixed rate of 11.5 percent. It should be obvious that the swaption is expiring with a positive value. MPK can enter into a swap to pay 11.5 percent fixed and receive LIBOR, whereas without the swap it would have to pay 12.75 percent fixed to receive LIBOR. Clearly the swap has value, but how much?

Suppose that MPK exercises the swaption, resulting in the entry into a swap to pay a fixed rate of 11.5 percent and receive LIBOR. Now suppose that it enters into the opposite swap in the market at the rate of 12.75 percent. Thus, it has the following positions:

A swap to pay 11.5 percent and receive LIBOR, plus

A swap to pay LIBOR and receive 12.75 percent.

The effect of LIBOR offsets on both swaps, leaving the following:

A position to pay 11.5 percent and receive 12.75 percent.

The net effect is that this position becomes a three-year annuity that pays 12.75% − 11.5% = 1.25%. It is a simple matter to determine the value of this annuity. The payments will be $10,000,000(0.0125) = $125,000. We find their present value using the discount factors for LIBOR obtained above:

$$\$125{,}000(0.8929 + 0.7901 + 0.6967) = \$297{,}463.$$

Thus, the swaption is worth $297,463 at expiration. In general, the payoff of a payer swaption at expiration is

[illegible]

where R is the swap rate at the swaption expiration, X is the exercise rate of the swaption, and the summation term captures the present value factors over the life of the swap. The payoff of a receiver swaption is

[illegible]

Returning to our example, MPK would exercise the swaption. MPK would then be engaged in a swap with a market value of $297,463. MPK could also choose to enter into a pay-floating, receive-fixed swap at the market rate of 12.75 percent. The net effect of these two swaps would be to create the annuity as described above, which has a market value of $297,463. Alternatively, MPK could just leave the swap created by exercise of the swaption in place. Its market value would be $297,463 instead of the normal zero market value when a swap is initiated. Had the swaption been structured to settle in cash, the short would have simply paid MPK $297,463. Of course, had the rate on swaps in the market been less than 11.5 percent, the swaption would have expired with no value. It should be apparent that when interest rates are high (at least high enough that the swap rate is above the swaption exercise rate), payer swaptions expire in-the-money and are, thus, somewhat like interest rate calls. When rates are low (at least low enough that the swap rate is below the swaption exercise rate), receiver swaptions expire in-the-money and are, thus, somewhat like interest rate puts.

Equivalence of Swaptions and Options on Bonds

Although swaptions may seem like complex instruments, it is actually quite easy to show that a swaption is identical to an option on a bond. Consider the MPK payer swaption. When it expires it will have the same value as that of a put option on a bond whose maturity corresponds to the maturity of the swap and whose coupon rate is the exercise rate of the swap. To simplify the problem, let us assume a notional principal of $1. Recall that we calculated its payoff value in the following manner:

$$\text{Max}\left(0, 0.1275 - 0.115\right)\left(\frac{360}{360}\right)\left(0.8929 + 0.7901 + 0.6967\right).$$

Recall that we found the value of 0.1275 in the following manner:

$$\left(\frac{1 - 0.6967}{0.8929 + 0.7901 + 0.6967}\right)\left(\frac{360}{360}\right) = 0.1275.$$

If we substitute the second equation into the first, we can write the payoff value of the swaption as

$$\text{Max}(0, 1 - 0.6967 - 0.115(0.8929 + 0.7901 + 0.6967)).$$

Now suppose that instead of the swaption, we purchase a put option with an exercise price of $1 on a bond that will have a three-year maturity when the option expires, annual interest payments of 11.5 percent, and a face value of $1. The value of this bond when the option expires would be the present value of three coupons of 0.115 and the present value of the principal of 1:

$$0.115(0.8929 + 0.7901 + 0.6967) + 1(0.8929).$$

The payoff of a put on this bond would be

$$\text{Max}(0,1 - 0.6967 - 0.115(0.8929 + 0.7901 + 0.6967)),$$

which is the same as the payoff of the payer swaption. Thus, a payer swaption is the same as a put option on a bond. Similarly, a receiver swaption can be shown to be equivalent to a call option on a bond.

Pricing Swaptions

Pricing swaptions is a somewhat advanced and complex topic that we cannot take up in any detail at this level. We can, however, appeal to the result of the previous section and gain some understanding of this process. As we showed, a swaption can be shown to be equivalent to an option on a bond. Thus, if we can price options on bonds, we can price swaptions. We have not taken up the pricing of options on bonds because this is a relatively advanced topic, requiring the construction of an arbitrage-free model of the evolution of the term structure. As an alternative, the Black-Scholes-Merton model can be adapted to pricing bonds, with considerable care and a number of caveats.

Forward Swaps

Throughout this book we have covered options and forward contracts. Since there are options on swaps, there will also be forward contracts on swaps. These instruments are called *forward swaps.* A forward swap commits the two parties to enter into a swap at a specific fixed rate. One party, the long, commits to enter into the swap to pay the fixed rate; the other commits to enter into the swap to receive the fixed rate. Because it is a forward contract, there is no cash flow up front.

Going back to our MPK example, suppose that instead of using a swaption, it wants to lock in the rate on the swap. Recall that when pricing a swap, we showed that a pay-fixed, receive-floating swap is equivalent to issuing a fixed-rate bond and using the proceeds to buy a floating-rate bond. In a similar manner, entering into a forward contract on a pay-fixed, receive-floating swap is equivalent to entering into a forward contract to issue a fixed-rate bond and buy a floating-rate bond. Entering into a forward contact to buy a floating-rate bond is trivial. Because the bond will be issued at par value, there is no uncertainty over its value when issued. Two parties simply agree that on a later date, one party will pay par value and receive a bond selling at par. Entering into a forward contract to issue a fixed-rate bond, however, does require solving for a given rate. The bond will need to be issued at par so that its value will equal that of the floating rate bond. Thus, its coupon will be adjusted so that it sells at par. When entering into a forward contract on this bond, the parties would need to solve for the appropriate coupon rate to make a forward contract on this bond be fairly priced.

This problem reduces to the simple problem of solving for the coupon rate on a par value bond to be issued at a later date. Going back to the MPK problem, suppose that, at the time the forward swap is created, MPK faces the following term structure:

Term	Rate	Discount Bond Price
360 days	9.00%	$B_0(360) = 1/(1 + 0.09(360/360)) = 0.9174$
720 days	10.06%	$B_0(720) = 1/(1 + 0.1006(720/360)) = 0.8325$
1080 days	11.03%	$B_0(1080) = 1/(1 + 0.1103(1080/360)) = 0.7514$
1440 days	12.00%	$B_0(1440) = 1/(1 + 0.12(1440/360)) = 0.6757$
1800 days	12.95%	$B_0(1800) = 1/(1 + 0.1295(1800/360)) = 0.6070$

MPK is interested in a forward contract expiring in two years on a swap that will last three years. To find the appropriate fixed rate, we shall need the forward rates two years ahead for periods of one, two, and three years. These will be found in the following manner.

The forward rates two years ahead are

$$\text{One year} = \left(\frac{1+0.1103\left(\frac{1080}{360}\right)}{1+0.1006\left(\frac{720}{360}\right)}-1\right)\left(\frac{360}{360}\right)=0.1080$$

$$\text{Two year} = \left(\frac{1+0.12\left(\frac{1440}{360}\right)}{1+0.1006\left(\frac{720}{360}\right)}-1\right)\left(\frac{360}{720}\right)=0.1161$$

$$\text{Three year} = \left(\frac{1+0.1295\left(\frac{1440}{360}\right)}{1+0.1006\left(\frac{720}{360}\right)}-1\right)\left(\frac{360}{1080}\right)=0.1238.$$

The discount factors using the forward rates would be[7]

$$B_0(720,1080)=\frac{1}{1+0.1080\left(\frac{360}{360}\right)}=0.9025$$

$$B_0(720,1440)=\frac{1}{1+0.1161\left(\frac{720}{360}\right)}=0.8116$$

$$B_0(720,1800)=\frac{1}{1+0.1238\left(\frac{1080}{360}\right)}=0.7292.$$

The rate on the forward swap would then be

$$\frac{1-0.7292}{0.9025+0.8116+0.7292}=0.1108.$$

Applications of Swaptions and Forward Swaps

The most obvious application of a swaption is when a company anticipates that it will need to enter into a swap at a later date and would like to establish a swap rate in advance. A swaption would permit the company to benefit from a favorable interest rate move, while protecting it against an unfavorable move. Of course, a swaption would require the payment of cash up front.

Consider a company that has already entered into a pay-fixed, receive-floating swap. If the company anticipates that it might want to terminate the swap, a receiver swaption would give it the right to do so by allowing it to enter into a pay-floating, receive-fixed swap at favorable terms. If the firm exercises the

[7]Here we must make a slight clarification in our notation. Using the first case as an example, the factor $B_0(720,1080)$ is the discount factor based on the forward rate established on day 0 that is based on a transaction to start on day 720 and end on day 1080.

swaption, it would have two swaps with equivalent streams of floating cash flows that are opposite in sign. It would also have two fixed streams at different rates. Of course, with both swaps in place, there would be some credit risk. Alternatively, if the counterparty to the swaption is the same counterparty to the original swap, then exercise of the swaption can be achieved by canceling the original swap and replacing it with either a lump sum cash payment of the market value or by leaving in place a stream of the net of the fixed rates on the two swaps. Since the firm might not know when it would want to terminate the swap, an American-style swaption would be preferred. In this manner, the company is using a swaption to give it flexibility to exit a swap.

Swaptions can also be used as a substitute for an option on a bond. When firms issue bonds and consider the possibility of wanting to change them from fixed rate to floating rate (or vice versa) at a later date, a swaption can provide that flexibility.

FIGURE [illegible] Illustration of Synthetic Fixed Rate Bond with Swaption

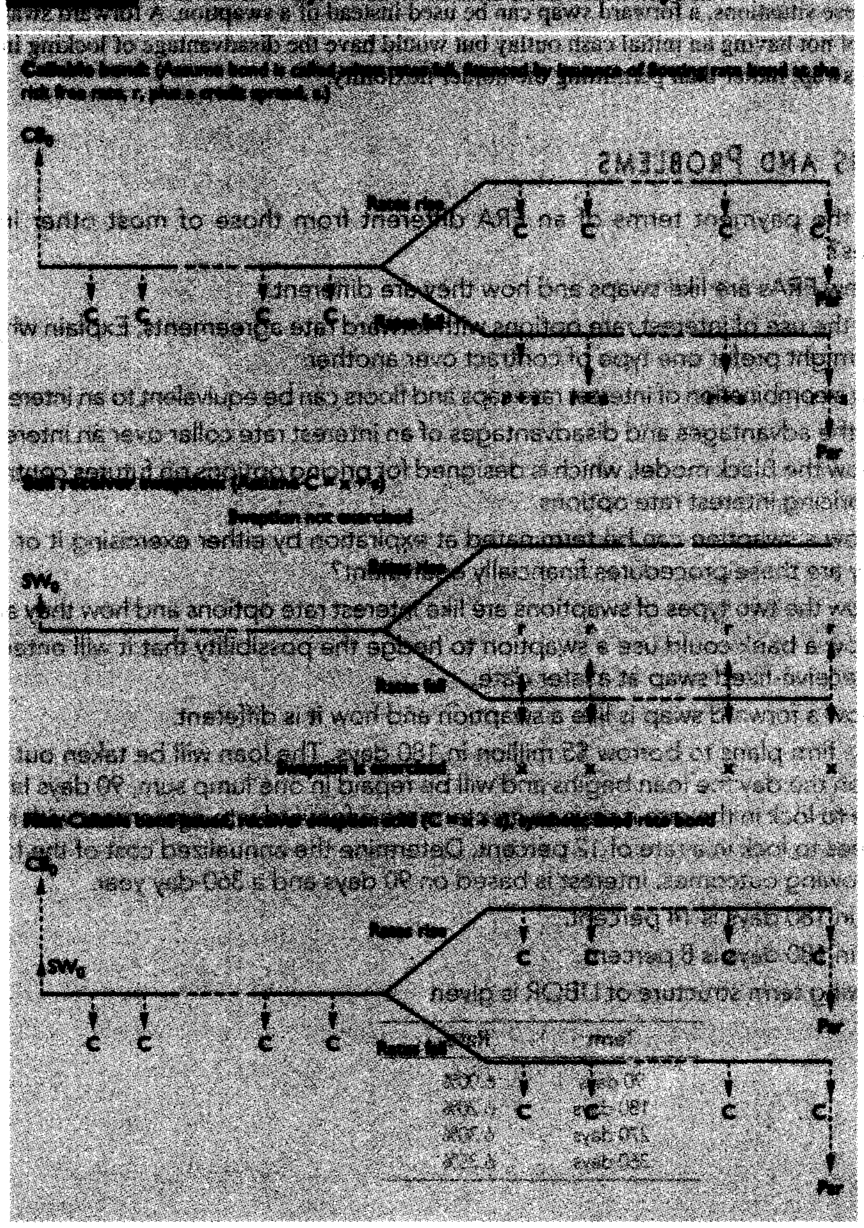

Another application of a swap is in creating synthetic callable debt. A callable bond is a bond in which the issuing firm has the right to call or retire it early. If a noncallable bond is issued but the firm would like to add the right to call it, the firm can purchase a receiver swaption. If properly structured, this swaption gives the firm the economic equivalent of an option to retire the bond early. Likewise, if the firm has issued a callable bond and no longer wants the right to call it, the firm can sell a receiver swaption, which effectively offsets the right to call the bond. Figure 13.5 presents a graphical depiction of a callable bond, a swaption, and the net position. The up arrows indicate receiving cash and the down arrows indicate paying cash.

A puttable bond is a bond which the holder of the bond has the right to sell back to the issuer early, thereby forcing the issuer to pay off the bond before its maturity. A bond that is not puttable can be effectively made puttable by selling a payer swaption. A bond that is puttable can be made nonputtable by buying a payer swaption.

In all of these situations, a forward swap can be used instead of a swaption. A forward swap would have the advantage of not having an initial cash outlay but would have the disadvantage of locking in the terms of the underlying swap, rather than permitting the holder flexibility.

QUESTIONS AND PROBLEMS

1. How are the payment terms of an FRA different from those of most other interest rate derivatives?
2. Explain how FRAs are like swaps and how they are different.
3. Compare the use of interest rate options with forward rate agreements. Explain why a financial manager might prefer one type of contract over another.
4. Show how a combination of interest rate caps and floors can be equivalent to an interest rate swap.
5. What are the advantages and disadvantages of an interest rate collar over an interest rate cap?
6. Explain how the Black model, which is designed for pricing options on futures contracts, can be used for pricing interest rate options.
7. Explain how a swaption can be terminated at expiration by either exercising it or settling it in cash. Why are these procedures financially equivalent?
8. Explain how the two types of swaptions are like interest rate options and how they are different.
9. Explain how a bank could use a swaption to hedge the possibility that it will enter into a pay-floating, receive-fixed swap at a later date.
10. Explain how a forward swap is like a swaption and how it is different.
11. Suppose a firm plans to borrow $5 million in 180 days. The loan will be taken out at whatever LIBOR is on the day the loan begins and will be repaid in one lump sum, 90 days later. The firm would like to lock in the rate it pays so it enters into a forward rate agreement with its bank. The bank agrees to lock in a rate of 12 percent. Determine the annualized cost of the loan for each of the following outcomes. Interest is based on 90 days and a 360-day year.
    a. LIBOR in 180 days is 14 percent.
    b. LIBOR in 180 days is 8 percent.
12. The following term structure of LIBOR is given

Term	Rate
90 days	6.00%
180 days	6.20%
270 days	6.30%
360 days	6.35%

a. Find the rate on a new 6 × 9 FRA.

b. Consider an FRA that was established previously at a rate of 5.2 percent with a notional principal of $30 million. The FRA expires in 180 days, and the underlying is 180-day LIBOR. Find the value of the FRA from the perspective of the party paying fixed and receiving floating as of the point in time at which the above term structure applies.

13. You are the treasurer of a firm that will need to borrow $10 million at LIBOR plus 2.5 points in 45 days. The loan will have a maturity of 180 days, at which time all the interest and principal will be repaid. The interest will be determined by LIBOR on the day the loan is taken out. To hedge the uncertainty of this future rate, you purchase a call on LIBOR with a strike of 9 percent for a premium of $32,000. Determine the amount you will pay back and the annualized cost of borrowing for LIBORs of 6 percent and 12 percent. Assume the payoff is based on 180 days and a 360-day year. The current LIBOR is 9 percent.

14. A large, multinational bank has committed to lend a firm $25 million in 30 days at LIBOR plus 100 bps. The loan will have a maturity of 90 days, at which time the principal and all interest will be repaid. The bank is concerned about falling interest rates and decides to buy a put on LIBOR with a strike of 9.5 percent and a premium of $60,000. Determine the annualized loan rate for LIBORs of 6.5 percent and 12.5 percent. Assume the payoff is based on 90 days and a 360-day year. The current LIBOR is 9.5 percent.

15. As the assistant treasurer of a large corporation, your job is to look for ways your company can lock in its cost of borrowing in the financial markets. The date is June 28. Your firm is taking out a loan of $20 million, with interest to be paid on September 28, December 31, March 31, and June 29. You will pay the LIBOR in effect at the beginning of the interest payment period. The current LIBOR is 10 percent. You recommend that the firm buy an interest rate cap with a strike of 10 percent and a premium of $70,000. Determine the cash flows over the life of this loan if LIBOR turns out to be 11 percent on September 28, 11.65 percent on December 31, and 12.04 percent on March 31. The payoff is based on the exact number of days and a 360-day year. If you have a financial calculator or a spreadsheet with an IRR function, solve for the internal rate of return and annualize it to determine the effective cost of borrowing.

16. You are a funds manager for a large bank. On April 15, your bank lends a corporation $35 million, with interest payments to be made on July 16, October 15, January 16, and next April 16. The amount of interest will be determined by LIBOR at the beginning of the interest payment period. On April 15, LIBOR is 8.0 percent. Your forecast is for declining interest rates, so you anticipate lower loan interest revenues. You decide to buy an interest rate floor with a strike set at 8 percent and a premium of $60,000. Determine the cash flows associated with the loan if LIBOR turns out to be 7.9 percent on July 16, 7.7 percent on October 15, and 8.1 percent next January 16. The payoff is based on the exact number of days and a 360-day year. If you have a financial calculator or spreadsheet with an IRR function, determine the internal rate of return and annualize it to determine your annualized return on the loan.

17. On January 15, a firm takes out a loan of $30 million, with interest payments to be made on April 16, July 15, October 14, and the following January 15, when the principal will be repaid. Interest will be paid at LIBOR based on the rate at the beginning of the interest payment period, using the exact number of days and a 360-day year. The firm wants to buy a cap with an exercise rate of 10 percent and a premium of $125,000 but is concerned about the cost. Its bank suggests that the firm sell a floor with an exercise rate of 9 percent for the same premium. The current LIBOR is 10 percent. Determine the firm's cash flows on the loan if LIBOR turns out to be 11.35 percent on April 16, 10.2 percent on July 15, and 8.86 percent on October 14. If you have a financial calculator or spreadsheet, determine the internal rate of return and annualize it to determine the cost of borrowing.

18. A bank is offering an interest rate call with an expiration of 45 days. The call pays off based on 180-day LIBOR. The volatility of forward rates is 17 percent. The 45-day forward rate for 180-day LIBOR is 0.1322 and the exercise rate is 12 percent. The risk-free rate for 45 days is 11.28 percent. All rates are continuously compounded. Use the Black model to determine how much the bank should receive for selling this call for every $1 million of notional principal.
19. A firm is interested in purchasing an interest rate cap from a bank. It has received an offer price from the bank but would like to determine if the price is fair. The cap will consist of two caplets, one expiring in 91 days and the other in 182 days. They will both have strikes of 7 percent. The forward rate applicable to the first caplet is 8 percent and the forward rate applicable to the second caplet is 8.2 percent. The 91-day risk-free rate is 7.1 percent and the 182-day risk-free rate is 7.3 percent. All rates are continuously compounded. The firm's best estimate of the volatility of forward rates is 16.6 percent. The notional principal is $10 million, and the payoff is based on 90-day LIBOR. Use the Black model to determine a fair price for the cap.
20. Consider a three-year receiver swaption with an exercise rate of 11.75 percent, in which the underlying swap is a $20 million notional principal four-year swap. The underlying rate is LIBOR. At the expiration of the swaption, the LIBOR rates are 10 percent (360 days), 10.5 percent (720 days), 10.9 percent (1,080 days), and 11.2 percent (1,440 days). Assume 360 days in a year. Determine the payoff value of the swaption.
21. A company wants to enter into a commitment to initiate a swap in 90 days. The swap would consist of four payments 90 days apart with the underlying being LIBOR. Use the term structure of LIBOR as given below to solve for the rate on this forward swap.

Term	Rate
90 days	10.2%
180 days	11.0%
270 days	11.6%
360 days	11.9%
450 days	12.2%

22. Suppose your firm had issued a 12 percent annual coupon, 15-year bond, callable at par at the 8th year. It is now two years later, so the bonds are not callable for another 6 years. At this time, new bonds could be issued at 8 percent, which is historically quite low, especially relative to the 12 percent coupon on the bond you issued two years ago. To provide a better matching of the interest-sensitivities of your assets and liabilities, you want to lengthen the duration of the bonds. How could you use swaptions to restructure the debt? Explain what happens assuming two subsequent future possibilities: rates going up and rates going down.
23. A firm has previously issued fixed rate non-callable debt. Because interest rates are perceived to be temporarily high, the firm would like to have the flexibility of calling the debt later when rates are expected to fall and replacing it with floating-rate debt. Explain how a firm can use swaptions to achieve this desired result. Also, identify and compare an alternative method that can be used to convert fixed-rate debt to floating-rate debt.
24. (Concept Problem) Use the Black model to determine a fair price for an interest rate put that expires in 74 days. The forward rate is 9.79 percent, and the exercise rate is 10 percent. The appropriate risk-free rate is 8.38 percent. All rates are continuously compounded. The volatility of forward rates is 14.65 percent. The put is based on $22 million notional principal and pays off based on 90-day LIBOR.

25. (Concept Problem) Consider a call option with an exercise rate of x on an interest rate, which we shall denote as simply L. The underlying rate is an m-day rate and pays off based on 360 days in a year. Now consider a put option on a $1 face value zero coupon bond that pays interest in the add-on manner (as in Eurodollars) based on the rate L. The exercise rate is X. Show that the interest rate call option with a notional principal of $1 provides the same payoffs as the interest rate put option if the notional principal on the put is $\$1(1 + x(m/360))$ and its exercise price, X, is $\$1/(1 + x(m/360))$. If these two options have the same payoffs, what does that tell us about how to price the options?

14

ADVANCED DERIVATIVES AND STRATEGIES

In this chapter we examine some advanced derivatives and strategies. Though it may seem like we are taking a leap forward, this is not necessarily the case. While some of these derivatives are complex, others are fairly simple and some are constructed by combining derivatives we have already covered.

We shall begin the chapter with a group called equity derivatives. Although these types of derivatives have been written primarily on stocks and stock portfolios, some of them are also written on bonds, interest rates, currencies, and commodities. The second group of derivatives we cover are structured notes and mortgage derivatives, which are related to interest rates. The third group is called *exotic* options. This is a term used to classify the most advanced types of options. Some of them hardly seem exotic, but they are typically referred to that way.

ADVANCED EQUITY DERIVATIVES AND STRATEGIES

An equity derivative is a derivative on a stock or stock index. While the over-the-counter market is much larger for interest rate derivatives, the over-the-counter equity derivatives market is growing rapidly. Recall that we have already studied at great length exchange-traded equity derivatives, such as options on stocks and stock indices, futures on stock indices, and options on stock index futures. While these contracts meet the needs of many investors, the specialization afforded by customized over-the-counter contracts can be increasingly worth the cost.

For example, many investors need derivatives that match specific portfolios. Exchange-traded derivatives are based on well-known indices, such as the S&P 500. To use a derivative on the S&P 500 to hedge a position in a portfolio that does not match the index induces some basis risk. While that risk may be acceptable for some investors, others would prefer to avoid it by using derivatives customized to their particular holdings. Such derivatives, called baskets, can be constructed by simply showing a derivatives dealer the composition of the portfolio. The dealer then constructs a derivative based on that specific portfolio.

Equity derivatives are especially useful in international investing. Equity derivatives on international stocks or indices allow an investor to capture the gains from investing in foreign securities without actually having to incur the additional transaction cost of acquiring those securities directly. In addition there are often costly foreign regulations to deal with and many countries impose a dividend withholding tax, which requires that some of the dividends be left in the foreign country to cover any future tax liabilities. Equity derivatives are often structured to pay off both capital gains and dividends, but since no money leaves the country, the dividend withholding tax is avoided. While investing directly in foreign securities imposes exchange rate risk, many equity derivatives are structured with a fixed exchange rate, which thereby avoids that risk.

Equity derivatives also permit fast, low-cost realignment of domestic portfolios. For example, an investor holding primarily large cap stocks and wanting to increase the allocation to small cap stocks can do so using derivatives on those two groups of stocks. This is much like our description of adjusting the risk of an equity portfolio in Chapter 11.

Naturally these features come at a cost that is incurred by the dealer and passed on to the end user. Fortunately the large size and sophisticated operations of most dealers result in efficiency and low costs, which benefit the end user.

Of course we have seen many of these benefits in previous chapters, but most of those instruments were exchange-traded equity derivatives. While we did cover equity swaps in Chapter 12, we take a look here at some specialized equity derivatives and strategies that meet needs not met by the instruments we have already covered.

Portfolio Insurance

The concept of insurance has been around for hundreds of years. Individuals and business firms routinely insure their lives and property against risk of loss. The concept extends easily to portfolios. In fact, we touched on the subject in Chapter 6 when we covered protective puts.

Suppose we own a portfolio consisting of N_S shares of stock and N_P puts. The stock price is S_0, the put price is P, the exercise price is X, the puts are European, and we assume no dividends on the stock. The value of the portfolio is

$$V = N_S S_0 + N_P P.$$

Letting $N_S = N_P$ and calling this N, we have

$$N = \frac{V}{S_0 + P}.$$

This tells us how many shares of stock and how many puts we can buy. At expiration the portfolio's value is

$$V_T = NS_T \qquad \text{if } S_T > X$$
$$V_T = NS_T + N(X - S_T) = NX \qquad \text{if } S_T \le X,$$

where S_T is the stock price when the put expires.

The worst possible outcome is that in which $S_T = 0$. Suppose we define V_{min} as the minimum value of V_T, which occurs when $S_T = 0$. Then $V_{min} = NX$ and, since N must also equal $V/(S_0 + P)$,

$$V_{min} = \frac{XV}{S_0 + P}.$$

This formula establishes the minimum value of the portfolio at expiration. Note that it occurs not only when $S_T = 0$ but also when $S_T \le X$.

Let us illustrate how this works. Suppose that on September 26 the market index is at 445.75 and the December 485 put option on the index is priced at $38.57. The option expires on December 19, which is 84 days away, so the time to expiration is 84/365 = 0.2301. The risk-free rate is 3.04 percent stated as a simple annual rate or 2.99 percent continuously compounded. The standard deviation is 15.5 percent.

Suppose we hold a diversified portfolio of stocks that replicates the index. The portfolio is worth \$44,575,000, which is equivalent to 100,000 units of the index. In other words we hold a portfolio that is weighted exactly like the index and is worth 100,000 times the index level.

The minimum level of the portfolio is

$$V_{min} = \frac{XV}{S_0 + P} = \frac{(485)(44,575,000)}{445.75 + 38.57} = 44,637,585.$$

Thus, the minimum level at which we can insure the portfolio is \$44,637,585. This means that if we own N shares and N puts, where

$$N = \frac{V}{S_0 + P} = \frac{44,575,000}{445.75 + 38.57} = 92,036,$$

the minimum value of the portfolio on December 19 is \$44,637,585. This is a guaranteed return of 0.0014 for 84 days, or $(1.0014)^{(365/84)} - 1 = 0.0061$ per year. This figure must be below the risk-free rate or an arbitrage opportunity would be possible. After all, how could we guarantee a minimum return on a risky portfolio greater than the risk-free rate?

We buy 92,036 shares and 92,036 puts. Suppose that at expiration the index is at 510:

$$\begin{aligned} \text{Value of stock} &= 92,036(\$510) = \$46,938,360 \\ + \text{ Value of puts} &= 92,036(\$0) \quad = \underline{\qquad\qquad 0} \\ \text{Total} &= \$46,938,360. \end{aligned}$$

This exceeds the minimum value.

If at expiration the index is at 450,

$$\text{Value of stock} = 92,036(\$485) = \$44,637,460 \text{ (by exercising the puts).}$$

While this amount appears to be slightly below the minimum, it is actually the same due to rounding off some of the previously computed values. The error is less than 1/1,000th of 1 percent.

Figure 14.1 shows the value of the insured stock-put portfolio when the put expires. The exact minimum cannot be read from the graph but is mathematically equal to \$44,637,585. The graph should look familiar. It is essentially the same as that of the protective put covered in Chapter 6. In this example, however, we are looking only at the value of the investor's position at expiration and not at the profit. The dotted line shows the value of the portfolio if it were uninsured.

Like any form of insurance, portfolio insurance entails a cost. By *cost* we do not necessarily mean commissions, bid-ask spreads, and so on. These certainly are important, but the cost of portfolio insurance is the difference in the return of the insured portfolio and the return of the uninsured portfolio when the market goes up. In other words, it is the return that is given up in bull markets, in which the insurance was not needed. This cost varies with the uninsured value of the portfolio. When the index ended up at 510, the insured portfolio was worth \$46,938,360. Had the portfolio not been insured, it would have consisted of 100,000 shares valued at \$510 each for a total value of \$51,000,000. The cost thus is \$4,061,640, or about 9.1 percent of the portfolio's original value. The difference between 100 percent and the cost of 9.1 percent, or 90.9 percent, is called the upside capture. It is the percentage of the uninsured return in a bull market that is earned by the insured portfolio. At low values of the portfolio, the cost of the insurance can be negative, which means that the insurance paid off more than it cost.

Because of put-call parity, an identical result is obtained with calls and risk-free debt, a strategy referred to as a fiduciary call. Let B be the price of a risk-free debt instrument and B_T be its face value when it matures.

FIGURE 14.1 Insured Portfolio: Stock–Put

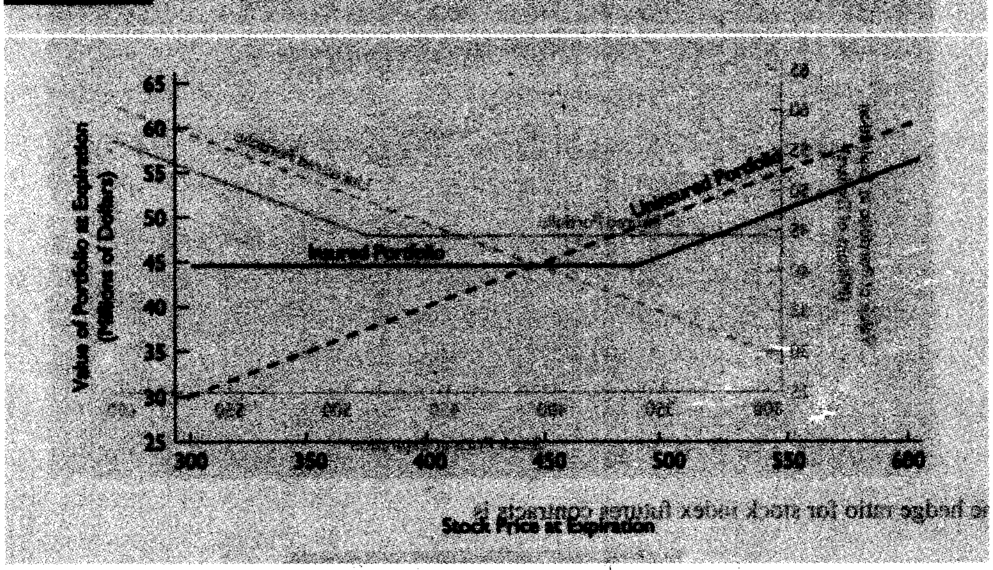

Suppose we own a portfolio of N_C calls and N_B debt. At expiration, the portfolio will be worth $V_T = N_B B_T$ if $S_T < X$ and $N_C(S_T - X) + N_B B_T$ if $S_T \geq X$. Again, the worst case is S_T going to zero. Call that result V_{min}, which will be the value of the debt at maturity, i.e., $V_{min} = N_B B_T$. Thus,

$$N_B = \frac{V}{S_0 B}$$

If we buy N_B debt, we can buy $N_C = (V - N_B B)/C$ calls. The number of calls can be shown to equal the number of shares and puts from the stock-put example. Thus, we also have

$$N_C = \frac{V}{S_0 + P}$$

Continuing with our numerical example, the call is worth \$2.65 and the debt has a face value of \$100. Thus, we would want \$44,637,585/\$100 = 446,376 of debt and 92,036 calls. At expiration if the index is at 510, the calls will expire \$25 in-the-money. With 92,036 calls worth \$25 and 446,376 debt worth \$100 each, the total portfolio value is \$46,938,500. If the index goes to 450, there will be 92,036 worthless calls and 446,376 debt worth \$100 each for a total of \$44,637,600, which is the same minimum in the stock-put case, subject to a round-off error. The example is illustrated in Figure 14.2.

A European put (or call) with the appropriate terms and conditions would obviously be a means of insuring a portfolio. While such options generally are not available on the options exchanges, it is possible to replicate the behavior of the stock-put or call–debt insured portfolio by continuously adjusting a portfolio of stocks and index futures or stocks and risk-free debt. This technique is referred to as dynamic hedging. When using stock index futures, dynamic hedging involves selling a number of futures contracts such that the portfolio responds to stock price movements the same way the stock-put or call–debt insured portfolio would respond.

FIGURE 14.2

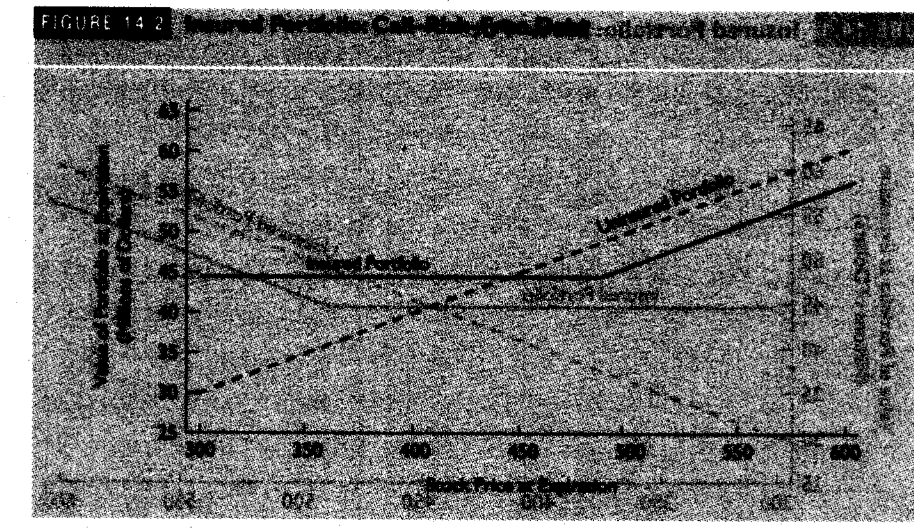

The hedge ratio for stock index futures contracts is

$$N_f = \left[\left(\frac{V_{min}}{X}\right)N(d_1) - \left(\frac{V}{S_f}\right)\right]e^{-r_c T}.$$

This gives the number of futures contracts. Since most futures contracts have a multiplier, this number should be divided by the multiplier. For example, if using the S&P 500 futures, the contract size is actually 250 times the quoted price. The above formula should then be divided by 250. Since futures require no outlay, the hedger can put all the funds in stock. The hedger must be sure to continuously adjust the number of futures contracts so that it always equals the value of N_f. Since the delta and T will change, this will necessitate frequent revisions.

An alternative approach to the use of futures in dynamic hedging is the use of risk-free debt. It is possible to combine stock and risk-free debt so that the portfolio behaves like the protective put or fiduciary call. The key assumption behind this result is that the risk-free debt price changes only as a result of time. It is not directly influenced by the stock price. This is a fairly reasonable assumption since the long-run correlation between stock returns and risk-free debt returns is nearly zero. The appropriate combination of stocks and risk-free debt results in the number of risk-free debt securities of:

$$N_B = \frac{V - N_s S}{B}.$$

The number of shares of stock is

$$N_s = \left(\frac{V_{min}}{X}\right)N(d_1).$$

As with the use of futures, these values will change continuously.

Both of these approaches work by equating the deltas of the dynamically hedged portfolios to the delta of the hypothetical portfolio of stock and puts. For example, for the December 485 call, the delta is 0.1574. This means that the put delta is $0.1574 - 1 = -0.8426$. A combination of one unit of stock and one put would have a delta of 1 (the stock's delta) plus -0.8426 (the put's delta), which equals 0.1574, the call's delta. When using stock index futures for dynamic hedging, the combination of all $44,575,000 in stock along with N_f futures has a delta of 0.1574, the same delta as would have been obtained if the hedger had owned 92,036 shares and 92,036 puts. When using the combination of stock and risk-free debt, the combination of N_S shares and N_B debt also has a delta of 0.1574.

Table 14.1 illustrates the appropriate calculations using stock and futures in a dynamic hedge, stock and risk-free debt in a dynamic hedge, and the exact hedge that would result if the necessary puts were available.[1]

Table 14.1 Dynamic Hedge Portfolio Insurance

Basic Information		Contract Information	
Current day	September 26	Futures price	448.83
Horizon day	December 19	Futures delta	1.0069
Time to horizon	0.2301	Futures multiplier	500
Stock price	445.75	Exercise price	485
Number of shares held	100,000	Put price	38.57
Market value of portfolio	$44,575,000	Call price	2.65
Risk-free rate	2.99%	Call delta	0.1574
Volatility of stock	15.5%	Put delta	−0.8426
Price of Risk-Free Debt	99.31		
Minimum portfolio value (V_{min})	$44,637,585		

Dynamic Hedge with Futures

Number of contracts needed:

$$N_f = \left[\left(\frac{44,637,585}{485}\right)(0.1574) - \left(\frac{44,575,000}{445.75}\right)\right]e^{-(0.0299)(0.2301)} / 500 \approx -170.$$

Dynamic Hedge with Risk-Free Debt

Number of shares held:

$$N_S = \left(\frac{44,637,585}{485}\right)(0.1574) \approx -14,486$$

Amount of debt needed:

$$N_B = \frac{44,575,000 - 14,486(445.75)}{99.31} \approx 383,827.$$

Effects of a $1 decrease in the stock price

New derivative prices:		
Put	$39.42	
Futures	447.82	
Stock plus put:		
Gain on stock	−$92,036	(92,036 shares × −$1)
Gain on put	78,230	(92,036 puts × ($39.42 − $38.57))
Net gain	−$13,806	(0.03% of value of portfolio)
Stock plus futures dynamic hedge:		
Gain on stock	−$100,000	(100,000 shares × −$1)
Gain on futures	85,850	(−170 contracts × (447.82 − 448.83) × 500)
Net gain	−14,150	
Stock plus debt dynamic hedge:		
Gain on stock	−$14,486	(14,486 shares × −$1)
Gain on debt	0	(debt price does not change due to stock price change)
Net gain	−$14,486	

[1]The old and new options and futures prices were calculated using the Black-Scholes-Merton model, as covered in Chapter 5, and the cost of carry model with continuous compounding, as covered in Chapter 9.

The lower panel shows that the differences in the three methods are very small. These differences can be explained by a combination of round-off error and the fact that a strategy equating deltas is only approximately correct if there is a non-infinitesimal change in the stock price, a point we introduced as the gamma effect in Chapter 5 and reexamine in Chapter 15. Moreover, dynamic hedging, or any delta-matching strategy, will require that the hedge ratio be adjusted continuously, which obviously cannot be done in the presence of transaction costs.

The dynamic hedging version of portfolio insurance was a widely used strategy in the mid-1980s. Its efficacy depends on the stock market having only small price changes within a short time interval. During the stock market crash of October 1987, however, stock prices experienced enormous jumps over short time intervals. While portfolios that were insured did better than those that were not, they failed to achieve their targets. The massive volume of selling done by portfolio insurers adjusting their hedge ratios was, probably wrongly, blamed for exacerbating the crash. Investors then became disillusioned with portfolio insurance. In response, however, it was not too long afterward that many new, customized equity products began to appear.

Equity Forwards

An equity forward contract is simply a forward contract on a stock, stock index, or portfolio. The principles are the same as those we studied previously in Chapter 9. An investor buying an equity forward simply enters into a contract with a counterparty, the seller, in which the buyer agrees to buy the stock, stock index, or portfolio from the seller at a future date at a price agreed upon today. In many cases, the equity instrument is a stock index or a portfolio of stocks. Rather than have the seller deliver stock to the buyer at expiration, the contract frequently specifies that it will be cash settled.

The price agreed upon by the two parties, called the forward price, was covered in Chapter 9. Recall that the price of a forward (or futures) contract on a stock is the current stock price, compounded at the risk-free rate minus the compound future value of the dividends. Since we have covered pricing at great length, we shall not repeat it here, but it may be useful for the reader to review the relevant parts of Chapter 9.

An interesting variation of a forward contract is the break forward. A break forward is a combination of spot and derivative positions that replicates the outcome of an ordinary call with one exception—the positions are structured such that the overall position costs nothing up front. This instrument is like a zero-cost call, except that it would be impossible to have an instrument that costs nothing up front and returns either zero or a positive amount, like an ordinary call. The break forward achieves this result by penalizing the investor if the option ends up out-of-the-money. An ordinary call pays off $\text{Max}(0, S_T - X)$. A break forward is a call that pays off if it expires in-the-money and incurs a charge if it expires out-of-the-money. Even the in-the-money payoff, however, can be negative.

To illustrate the break forward and to keep the explanation simple, we assume no dividends on the stock. A break forward will pay off $S_T - K$ if $S_T > F$ and $F - K$ if $S_T \leq F$. The value K is the sum of the compound future value of an ordinary call on the stock with exercise price F plus the compound future value of the stock. The latter term, you know, is the forward price. Thus, K will exceed F so when $S_T \leq F$, the payoff, $F - K$, is definitely negative. It is also possible that $S_T - K$ can be negative. If the break forward were structured as a put, the exercise price would be the forward price plus the compound value of the put price.

A break forward contract is identical to an ordinary long call with an exercise price of F and a loan in which the investor receives the present value of $K - F$ and promises to pay back $K - F$. Table 14.2 illustrates this result. Recall from Chapters 2–7 that we examined some results for DCRB options. We shall use the same DCRB options for many of the examples in this chapter.[2] The value of the break forward at expiration is illustrated in Figure 14.3.

[2]In practice, many of these types of contracts are based on the S&P 500 or a specific portfolio. We use the DCRB examples for convenience and because you should by now be quite familiar with these options.

DERIVATIVES TOOLS

Concepts, Applications, Extensions
Portfolio Insurance in a Crashing Market

On October19, 1987 the Dow Jones Industrial Average fell over 22 percent in one day. Many factors were blamed for the crash but none took more criticism than portfolio insurance.

As discussed in this chapter, portfolio insurance is a strategy in which a portfolio is protected by either buying puts, combining calls with risk-free bonds, or replicating a protective put using a dynamic hedging strategy. Exchange-listed puts with the necessary terms are not available for the customized portfolios of many institutional investors; hence, portfolio insurance is typically implemented using dynamic hedging.

As we have discussed in this chapter, dynamic hedging attempts to replicate a position combining an asset with a put. Futures contracts can be combined with the portfolio to obtain the same sensitivity of the portfolio combined with a put. Because it is based on calculus, the formula is correct only for a small change in the underlying. Thus, portfolio insurance works well only when the market makes small moves and only if the market is liquid.

One of the primary vendors of portfolio insurance was the California-based Leland O'Brien Rubinstein, known as LOR, which had been formed in 1982. Hayne Leland and Mark Rubinstein are finance professors at the University of California at Berkeley who had developed the theory underlying portfolio insurance. John O'Brien, an investment practitioner, brought marketing skills to the firm. LOR successfully implemented this strategy for numerous clients. Because the product was simple and required no specialized trade secrets, it was quickly copied. Soon portfolio insurance strategies were associated with billions of dollars of money in the U.S. stock market.

But on October 19, 1987 the Dow dropped about 200 points in the first two hours and about 200 points in the last two hours. Although the Dow made a 100-point move upward during mid-day, it was largely a day of declines. The Dow was at it highest at the opening and its lowest at the close. It was a day unheard of in stock market history.

Following large market movements, portfolio insurers must make rapid adjustments to their positions. Consider a portfolio insured with a specific number of short positions in futures. If the market makes a large downward move, the portfolio would then require more short futures contracts. The portfolio insurer then attempts to sell more futures contracts, and these transactions must be executed before another large move. If the portfolio insurer needs to sell a large number of contracts, the market could mistakenly infer that the portfolio insurer believes the market is going down further, a negative but misleading signal. The transaction is not based on a belief that the market is going down further. Even if the portfolio insurer believes the market is going up, the same transaction—selling more futures contracts—would be required. In addition, there is the possibility that the large number of futures contracts would strain the capacity of the market and liquidity would suffer. Prices would rapidly fall further. Portfolio insurance models would then dictate further selling, leading to an endless spiral. Of course if prices rose rapidly, portfolio insurance models would dictate rapid buying, but this was not the case that day. To make matters worse, many markets were shut down at various times, effectively prohibiting anyone from getting in or out.

At the end of the day, portfolio insurance took considerable blame for the bloodbath. In fact, the portfolio insurance strategy and many portfolio insurance firms would never be the same. This criticism was unjust. While it is probably true that portfolio insurance strategies were responsible for some of the decline, much of the problem lay in the fact that the market panicked, misinterpreted signals, and then shut itself down.

In fact portfolio insurance resulted in significant savings to investors who had it. It did not live up to its promises, however, because it could not operate well in rapidly moving markets. Clients were dissatisfied and terminated their relationships with portfolio insurance firms. Most (including LOR) did not survive. Since that time, the market has developed a variety of alternative and successful derivative instruments and strategies. For that, we can thank portfolio insurance and companies like LOR. Financial products that succeed could never be developed without knowing that products like portfolio insurance worked well in some circumstances and badly in others.

Table 14.2 Payoffs from Break Forward

		Payoffs from Portfolio Given Stock Price at Expiration	
Instrument	**Current Value**	$S_T \le F$	$S_T > F$
Long Call	$C_e(S_0,T,F)$	0	$S_T - F$
Loan	$-(K - F)e^{-r_cT}$	$-(K - F)$	$-(K - F)$
	0	$F - K$	$S_T - K$

The example is based on a stock price of 125.94, a time to expiration of 0.0959 (35 days/365), a risk-free rate of 4.46 percent continuously compounded, a volatility of 83 percent, and 100 units of the break forward. The forward price of DCRB is $125.94e^{0.0446(0.0959)} = 126.48$. This value is the exercise price of the break forward. As noted above, the break forward is equivalent to a loan and a call with exercise price of F, so we need the price of a call expiring in 35 days with an exercise price of 126.48. Using the Black-Scholes-Merton model, that call price would be 12.90. Next we need to know the value of K. This will equal the forward price plus the compound future value of the call price. This is $126.48 - 12.88e^{0.0446(0.0959)} = 139.41$. Thus, the break forward will be worth $126.48 - 139.41 = -12.93$ if $S_T \le 126.48$ and $S_T - 139.41$ if $S_T > 126.48$. In the figure all values are multiplied by $100.

FIGURE 14.3 Break Forward

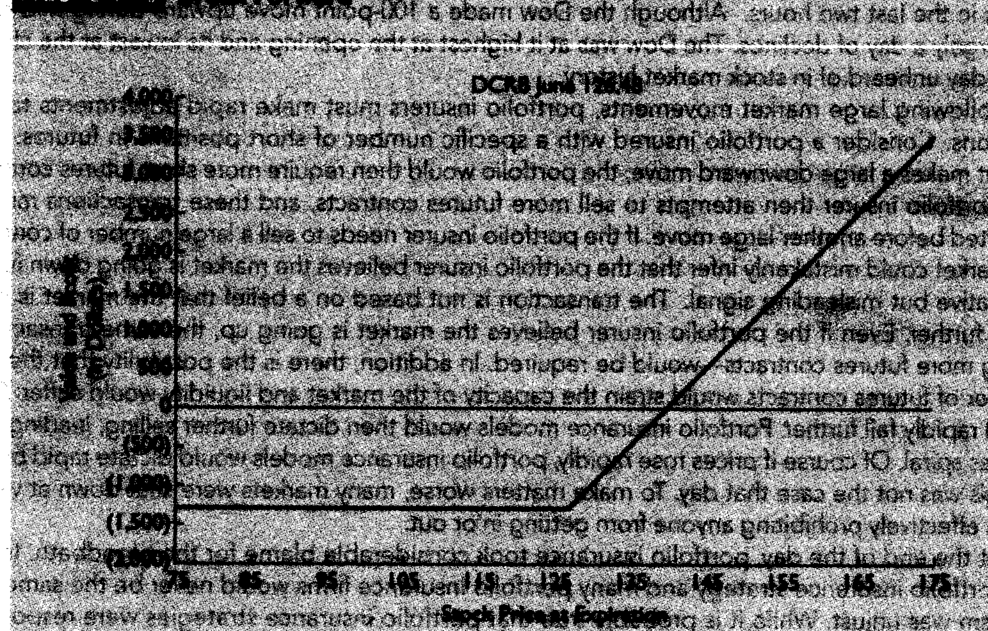

Since the up-front cost is zero, the value at expiration is also the profit. While this figure looks exactly like a call, it is important to remember that a break forward incurs a payment if it expires out-of-the-money. This is the penalty for having to pay nothing for it up front.

It should be apparent that a break forward is similar to a forward contract and a call option. Like a forward contract, it requires no initial outlay but can have negative value at expiration. Like a call, it has a limited loss. Due to the parity relationships that exist between puts, calls, and forwards, a break forward can also be constructed by entering into a long forward contract at the price F, borrowing the present value of K − F and buying a put with an exercise price of F.

Like a forward contract or a swap, a break forward has a value of zero at the start. Later during its life, the value will change to positive or negative. Thus, we need to know how to determine its value during its life. This is a simple procedure. Since we know that a break forward is like an ordinary call and a loan, we can value it by valuing the ordinary call and the loan. Thus, suppose we are at time t prior to expiration. The call is worth $C_e(S_t, T - t, F)$, and the loan is worth $-(K - F)e^{-r_c(T-t)}$. The latter is the loan payoff, the amount that would be owed if the loan were repaid now. In the example we worked here, suppose we are 15 days into the life of the break forward, and DCRB is at 115.75. The other inputs are F = 126.48 (the exercise price), $r_c = 0.0446$, $K = 139.41$, $\sigma = 0.83$, $T - t = 20/365 = 0.0548$. Using a spreadsheet for greater precision, we obtain a Black-Scholes-Merton value for the call of 5.05. The loan value is $-(139.41 - 126.48)e^{-0.0446(0.0548)} = 12.90$. Thus, the value of the break forward is 5.05 − 12.90 = −7.85. Since the break forward can have a negative payoff at expiration, it can obviously have a negative value during its life.

A more general version of a break forward is a pay-later option. In this case, the buyer of the option simply borrows the premium. Using a European call as an example, this is an ordinary call with a premium of $C_e(S_0, T, X)$ that also includes a loan in the amount of $C_e(S_0, T, X)$. At expiration, the borrower either exercises the call or not but is required to pay back the amount $C_e(S_0, T, X)e^{r_cT}$. This instrument is, in effect, an option in which the premium does not have to be paid until expiration. With an option of this type, any exercise price can be chosen. A break forward is a special case in which the exercise price is the forward price. Later in this chapter, we look at an option in which the premium is not paid until expiration and is paid only if the option expires in-the-money.

Equity Warrants

Warrants have been around much longer than exchange-listed options. A warrant is an option written by a firm on its own stock and usually offered with a bond issue. Anyone purchasing the bond receives one or more warrants, which give the bondholders the right to buy the stock later. The original life of a warrant is usually three to ten years. Warrants can be priced similarly to ordinary options except that their exercise dilutes the value of the stock and this must be taken into account. Many warrants trade on stock exchanges.

From time to time, warrants are written by a financial institution, offered for sale to the public, and typically based on a stock index. Many of these warrants on foreign indices offer the interesting feature that even though the underlying index is stated in units of the foreign currency, the payoff is fixed in U.S. dollars. For example, suppose there is a call warrant on the Japanese index, the Nikkei 225 with a fixed exchange rate of ¥100 per dollar. Let the exercise price be ¥19,000. At expiration let the Nikkei be ¥19,950. Thus, the call warrant expires in-the-money by ¥950.

Because the index value is based on prices quoted in yen, the warrant ends up worth ¥950. This figure is automatically converted to dollars at ¥100 per dollar, so the warrant pays off $9.50. This type of fixed exchange rate derivative is called a quanto.

Quantos are particularly attractive because they permit investors to earn returns in foreign markets without exchange rate risk.

Equity-Linked Debt

Equity-linked debt is a combination of a call option and a bond that guarantees a return of principal and either no interest, a small amount of interest, or a guarantee of returning a given fraction of the principal. In addition, it provides an extra payment based on how a stock or stock index performed. Instruments of this sort were first offered in 1987 in the form of bank certificates of deposit and were available to the general public. These original market-indexed CDs did not prove to be very popular, and they were abandoned. In the last few years they have resurfaced with some alterations.

Suppose a bank or investment banking firm makes the following offer: Purchase a one-year zero coupon bond paying 1 percent interest and receive 50 percent of any upside gain on the S&P 500. Each unit is sold with a principal amount of $10. Currently a one-year zero coupon bond without the S&P feature would pay 5 percent compounded annually. The S&P 500 is at 1500, its standard deviation is 12 percent, and its dividend yield is 1.5 percent. Is this a good deal?

If you invest $10, you are guaranteed to receive $10(1.01) = $10.10. Since the opportunity cost is 5 percent, the present value of this amount is $10.10/1.05 = $9.62. This means that you are implicitly paying $10.00 − $9.62 = $0.38 for 50 percent of the S&P 500 upside gain. Obviously the upside gain feature is like a call option. We can easily determine its value.

Since the option payoff cannot be earned early, it is a European call and the Black-Scholes-Merton model applies with a slight twist. The payoff is $\$10(0.5)\text{Max}(0,(S_T - 1500)/1500)$ where S_T is the S&P 500 at expiration. This can be written as $(5/1500)\text{Max}(0,S_T - 1500)$, making it equivalent to 5/1500th of an at-the-money call. The option value will, therefore, be 5/1500th of the Black-Scholes-Merton value. The inputs to the Black-Scholes-Merton calculation are $S = 1500$, $X = 1500$, $r = \ln(1.05) = 0.0488$, $T = 1$, $\sigma = 0.12$, and $\delta = 0.015$. Plugging into the Black-Scholes-Merton model gives $C = \$96.81$. Thus, the option value to the holder of this security is (5/1500)($96.81) = $0.32. This is slightly less than the implied cost of the option of $0.38, which is not surprising given that the firm offering the deal must cover its costs and make a profit.

These equity-linked securities sometimes have other features. Some pay off based on the average price over the last 10 days or so before expiration. Others pay no interest or even implicitly a negative interest rate by promising to return a given percentage but less than 100 percent of the amount invested. The holder then gets a greater percentage of the gain in the S&P 500. Some of the securities make their minimum interest payments quarterly or semiannually. Others pay off based on any depreciation in the index, which makes the option feature like a put.

These securities have at times been created by stock and options exchanges but mostly trade on the over-the-counter market.

ADVANCED INTEREST RATE DERIVATIVES

In Chapter 12 we covered interest rate swaps. In Chapter 13 we encountered some other derivatives on interest rates. These consisted of forward rate agreements, options, and swaptions. In this section we look at some other widely used interest rate instruments that are not technically derivatives but are generally considered as such.

Structured Notes

Corporations routinely issue debt obligations with maturities in the range of 2–10 years. These instruments are usually called notes. Sometimes the notes carry fixed rates. Other times the notes have variable rates, such as those tied to LIBOR. In the early 1990s many corporations began issuing notes with derivative transactions attached so as to change the return pattern from the standard fixed or LIBOR-floating type. These instruments have come to be known as structured notes.

Most structured notes are issued by corporations of extremely high credit quality. Consequently, there is very little credit risk. The notes are typically designed with a specific user in mind. The user wants a particular exposure and normally plans to hold the notes until maturity. Thus, these tend to be fairly illiquid instruments.

Some of these notes have coupons indexed to the CMT rate. Recall that we discussed the CMT or constant maturity treasury rate in Chapter 12. It is the rate on a government security of a fixed maturity. If there is no government bond outstanding with the desired maturity, the CMT rate is obtained by interpolating the rates on bonds with surrounding maturities. For example, the coupon could be the CMT rate plus 50 basis points. Some structured notes have leveraged coupons, meaning that, for example, the coupon might be 1.5 times the CMT rate. Let us say the CMT rate was currently 5.5 percent. Then the coupon would be $5.5(1.5) = 8.25$ percent. Suppose that the CMT rate went up 100 basis points to 6.5. Then the structured note coupon would go up to $6.5(1.5) = 9.75$, a 150-basis-point increase.

One interesting type of structured note is the range floater. This instrument pays interest only if a specified interest rate, such as LIBOR, stays within a given range. The ranges are sometimes different for each year of the note's life. As an example, consider a two-year note in which the coupon is LIBOR + 3 percent provided that:

For the first year, LIBOR is between 0 and 6 percent, and
For the second year, LIBOR is between 0 and 7 percent.

In other words, if LIBOR is outside the specified range, the coupon is zero. Otherwise, the coupon is LIBOR + 3 percent, which would be considered above normal for a two-year note. You can see that this is a bet on LIBOR staying within a range. To the investor this transaction would be like a standard floating-rate note and the sale of an interest rate cap. If rates rise above a certain threshold, the investor gives up all the interest.

A popular and controversial type of structured note is the inverse floater, sometimes called the reverse floater. Whereas a standard floating-rate note pays interest at a rate that changes directly with interest rates in the market, an inverse floater pays interest that changes in an opposite manner to rates in the market. In other words, as interest rates go up, the coupon on an inverse floater goes down. This is usually established by setting the coupon at a formula like Y − LIBOR, where Y is a specified fixed rate. Sometimes it is simply a multiple of the fixed rate on a plain vanilla swap.

Suppose Y is 12 percent and LIBOR is currently 5.5 percent. Then the coupon will be 6.5 percent. If LIBOR moves up to 7 percent, the coupon will be 5 percent. If LIBOR moves down to 3 percent, the coupon will be 9 percent. It should be obvious that if LIBOR exceeds 12 percent, the coupon is negative, meaning that the lender rather than the borrower would pay the interest! In practice, however, most inverse floaters specify that the minimum interest rate will be zero or perhaps a very small positive rate.

The issuer of an inverse floater could hedge its position with a plain vanilla swap paying LIBOR and receiving a fixed rate. For the case of LIBOR < 12 percent, the issuer's payment would be 12 minus LIBOR on the structured note and + LIBOR minus the fixed rate on the swap for a total payment of 12 minus the fixed rate. The issuer would have created a synthetic fixed-rate loan at this rate. If LIBOR exceeds 12 percent, the issuer's payment would be the same as a pay floating-receive fixed swap. If the issuer does not want such a position, it might wish to buy a cap with a strike rate of 12, which would pay it LIBOR minus 12 if LIBOR exceeds 12 percent. In that case, for LIBOR of more than 12, the issuer's total payment would be −(LIBOR −12) from the cap and LIBOR minus the fixed rate on the swap for a total payment of 12 minus the fixed rate. In cases where the inverse floater payment is twice the fixed rate minus LIBOR, the net payment of the issuer after the swap will be twice the fixed rate minus the fixed rate or simply once the fixed rate. In that case, the issuer has offered an inverse floater but effectively pays only the fixed rate.

Some inverse floaters are leveraged with formulas such as three times the fixed-rate minus twice LIBOR. Many inverse floaters have caps on the maximum rate. Some are de-levered by using formulas like Y − 0.75LIBOR, where the floating coupon moves less than the movement in LIBOR.

There are many variations found in structured notes. Some have coupons determined by the difference between two interest rates, such as the TED spread, and others have coupons tied to foreign interest rates but with payments made in the domestic currency. Structured notes have also been based on complex formulas reflecting differences in rates and bond prices at different ends of the term structure.

The success of structured notes is attributed to the fact that they offer return possibilities tailored to specific investors. Unfortunately, structured notes have also been aggressively marketed in some cases to investors who were not well suited for them. Large interest rate changes coupled with high leverage resulted in large losses for investors who were not aware they had so much exposure.

Mortgage-Backed Securities

The fixed-rate home mortgage is the primary means by which individuals purchase houses. Most homeowners pay off their mortgages over a period of 15 to 30 years with fixed and equal monthly payments, consisting of interest on the remaining balance and a contribution to the principal. Nearly every mortgage also provides the homeowner the right to prepay the mortgage, which will frequently be done if interest rates decline sufficiently to cover the transaction cost of refinancing. A certain number of prepayments occur due to demographic factors, such as a family moving to another city. Mortgage lenders are primarily banks, savings and loan associations, and mortgage companies. These financial institutions that originate mortgage loans are exposed to the risk of interest rate changes. If interest rates increase, the value of a mortgage decreases. If interest rates decrease, the value of a mortgage increases but the increase is limited by the fact that lower interest rates encourage homeowners to refinance. If that happens, the mortgage holder receives the principal early and is forced to reinvest it at a lower rate. This is referred to as prepayment risk.

Many mortgages are combined into portfolios; claims on the portfolio are sold to investors in the form of mortgage-backed securities. The process of combining loans into portfolios and selling claims on the portfolio is called securitization. Not all securitized portfolios are mortgages. Credit card receipts and other types of loans are also securitized, but the mortgage securitization market is one of the largest and most complex. The process of securitization provides depth to the market because lenders know that they can sell the loans they hold and receive immediate cash. This increases the number of willing lenders and provides individual and institutional investors access to investment opportunities in markets that would otherwise be accessible only to financial institutions. Securitized mortgages are nearly always guaranteed against default by the purchase of credit insurance, often from a federal agency such as the Government National Mortgage Association (GNMA).

Securities that are issued based on mortgage portfolios are generally considered to be derivatives because they derive their values from the values of the underlying mortgages. They are, however, somewhat different from traditional derivatives. Though most mortgages have imbedded prepayment options, mortgage securities are more like mutual funds. They are investments in a security that represents holdings of another security. Nonetheless, we shall cover them here because it is important to know something about them and in practice, nearly everyone, rightly or wrongly, thinks of them as derivatives.

Mortgage Pass-Throughs and Strips Portfolios of mortgages on which claims are sold are called mortgage pass-throughs. An investor who purchases a mortgage pass-through receives a monthly check consisting of a proportional share of interest and principal on the portfolio of mortgages, less a servicing fee. Mortgage pass-throughs, however, are still subject to the risk of interest rate changes, which, as we noted above, result in changes to the values of the securities and affect the frequency with which homeowners refinance their mortgages. Consequently, mortgage pass-through values can change dramatically with changes in interest rates.

Some mortgage pass-throughs are decomposed into mortgage strips, each consisting of an interest or principal stream, which are sold separately. An investor can purchase either or both strips. The interest-only strip, called an IO, provides the investor with the interest payments on the portfolio of underlying mortgages. The principal-only strip, called a PO, provides the investor with the principal payments on the portfolio of underlying mortgages. If interest rates decrease and prepayments accelerate, the holder of an IO will find the value of his position decreasing, perhaps quite dramatically. Thus, it sounds as if the holder of an IO would prefer for interest rates to increase. That will be the case to a limited extent. Beyond a point, however, increases in interest rates will reduce the present value of the interest-only stream.

Understanding the mechanics of mortgage-backed securities is difficult. At this introductory level, our objective is to just grasp the fundamentals. We shall do this by looking at a simple mortgage portfolio, one consisting of a single mortgage. Let that mortgage be a $100,000 loan at 9.75 percent for 30 years. Although most mortgages have monthly payments, we shall use annual payments in this example. This enables us to show the entire amortization schedule, which is in Table 14.3.

First let us determine the annual payment, a simple problem in the mathematics of compound interest. The payment will equal the loan amount divided by the discount factor for an annuity:

$$\text{Payment} = \frac{\$100,00}{(1-(1.0975)^{-30})/0.0975} = \$10,387.$$

Each payment consists of interest, which is 0.0975 times the balance owed, plus the contribution to the principal, which is $10,387 minus the interest. For example, the 10th payment of $10,387 consists of interest of $91,436(0.0975) = $8,915 and the contribution to principal of $1,472, which leaves a balance of $91,436 – $1,472 = $89,964.[3]

Table 14.3 Amortization Schedule for a 30-Year, $100,000 Mortgage at 9.75 Percent

		Annual Payments			
Payment No.	Balance before Payment	Payment	Interest	Principal	Balance after Payment
1	$100,000	$10,387	$9,750	$637	$99,363
2	99,363	10,387	9,688	699	98,663
3	98,663	10,387	9,620	768	97,896
4	97,896	10,387	9,545	843	97,053
5	97,053	10,387	9,463	925	96,128
6	96,128	10,387	9,373	1,015	95,113
7	95,113	10,387	9,274	1,114	94,000
8	94,000	10,387	9,165	1,222	92,777
9	92,777	10,387	9,046	1,342	91,436
10	91,436	10,387	8,915	1,472	89,963
11	89,963	10,387	8,771	1,616	88,348
12	88,348	10,387	8,614	1,773	86,574
13	86,574	10,387	8,441	1,946	84,628
14	84,628	10,387	8,251	2,136	82,492
15	82,494	10,387	8,043	2,344	80,147
16	80,147	10,387	7,814	2,573	77,574
17	77,574	10,387	7,563	2,824	74,750
18	74,750	10,387	7,288	3,099	71,651
19	71,651	10,387	6,986	3,401	68,250
20	68,250	10,387	6,654	3,733	64,517
21	64,517	10,387	6,290	4,097	60,420
22	60,420	10,387	5,891	4,496	55,924
23	55,924	10,387	5,453	4,935	50,989
24	50,989	10,387	4,971	5,416	45,573
25	45,573	10,387	4,443	5,944	39,629
26	39,629	10,387	3,864	6,524	33,105
27	33,105	10,387	3,228	7,160	25,946
28	25,946	10,387	2,530	7,858	18,088
29	18,088	10,387	1,764	8,624	9,465
30	9,465	10,387	923	9,465	0

Note: Annual payments are used as an approximation. Most mortgages have monthly payments.

[3]The figures in the table reflect more precise calculations so they occasionally differ slightly from manually calculated figures.

Now let us examine the values of IOs and POs. The value of the IO strip will be the present value of the column of numbers representing the interest. The value of the PO strip will be the present value of the column of numbers representing the principal.

To determine these values we need two pieces of information. One is the rate at which to discount the cash flows. Assuming the payments are guaranteed by a federal agency, we can discount these at a rate appropriate for a Treasury security. So at the present, let us assume a rate of 7 percent. We also have to make an assumption about when the loan will be paid off. Although it is a 30-year loan, it has a good chance of being paid off early. We shall assume that it is paid off in its twelfth year, which is considered to be an average payoff date for 30-year mortgages.

At the current interest rate of 7 percent with the mortgage prepaid in the twelfth year, the value of the IO strip is found by discounting the 12 interest payments at 7 percent:

Value of IO strip at 7 percent interest and prepayment in year 12
$$= \$9{,}750(1.07)^{-1} + \$9{,}688(1.07)^{-2} + \cdots + \$8{,}614(1.07)^{-12} = \$74{,}254.$$

The value of the PO strip is found by discounting the 12 principal payments, taking into account that the final principal payment will include the full payoff of $86,574:

Value of PO strip at 7 percent interest and prepayment in year 12
$$= \$637(1.07)^{-1} + \$699(1.07)^{-2} + \cdots + (\$1{,}773 + \$86{,}574)\,(1.07)^{-12} = \$46{,}690.$$

The overall pass-through would be worth the sum of these two or

$$\text{Value of pass-through} = \$74{,}254 + \$46{,}690 = \$120{,}944.$$

Now assume the interest rate falls to 6 percent and assume the homeowner prepays in two years. Then we have

Value of IO strip at 6 percent interest and prepayment in year 2
$$= \$9{,}750(1.06)^{-1} + \$9{,}688(1.06)^{-2} = \$17{,}820,$$

which is a loss of 76 percent.

Value of PO strip at 6 percent interest and prepayment in year 2
$$= \$637(1.06)^{-1} + (\$699 + \$98{,}663)(1.06)^{-2} = \$89{,}034,$$

which is a gain of 91 percent. The overall pass-through is worth

$$\text{Value of pass-through} = \$17{,}820 + \$89{,}034 = \$106{,}854,$$

a loss of almost 12 percent.

Assume that the interest rate is at 8 percent and prepayment still occurs in year 12. Then we have:

Value of IO strip at 8 percent interest and prepayment in year 12
$$= \$9{,}750(1.08)^{-1} + \$9{,}688(1.08)^{-2} + \cdots + \$8{,}614(1.08)^{-12} = \$70{,}532,$$

a loss of 5 percent.

Value of PO strip at 8 percent interest and prepayment in year 12
$$= \$637(1.08)^{-1} + \$699(1.08)^{-2} + \cdots + (\$1{,}773 + \$86{,}574)(1.08)^{-12} = \$42{,}128,$$

a loss of almost 10 percent. The overall pass-through is worth

$$\text{Value of pass-through} = \$70{,}532 + \$42{,}128 = \$112{,}660,$$

a loss of almost 7 percent.

It would appear that IO holders lose either way, but that is not exactly how it works in practice. IO and PO strips represent portfolios of mortgages. Each mortgage will not prepay at the same time. Any increase in interest rates is likely to push the average prepayment date further out. If in our example, the interest rate change to 8 percent pushes the prepayment to year 14, we have

Value of IO strip at 8 percent interest and prepayment in year 14

$$= \$9{,}750(1.08)^{-1} + \$9{,}688(1.08)^{-2} + \cdots + \$8{,}251(1.08)^{-14} = \$76{,}445,$$

which is a gain of about 3 percent.

Value of PO strip at 8 percent interest and prepayment in year 14

$$= \$637(1.08)^{-1} + \$699(1.08)^{-2} + \cdots + (\$2{,}136 + \$82{,}492)(1.08)^{-14} = \$37{,}276,$$

which is a loss of about 20 percent. The overall pass-through is worth

$$\text{Value of pass-through} = \$76{,}445 + \$37{,}276 = \$113{,}721,$$

a loss of about 6 percent.

A mortgage pass-through decreases in value with increases in interest rates. When rates decrease, however, the gains on the pass-through are limited because of the increased prepayment rate. The pass-through is not as volatile as the IO or PO individually but it can still be fairly volatile and has limited gains because of the prepayment risk.[4]

The main point from this demonstration is that mortgage securities are complex and that POs and IOs can be extremely volatile. Not only are they affected by the discounting effect of interest rates but they are extremely sensitive to the effect that interest rate changes have on prepayments. As you might well imagine, it is difficult to predict prepayments, making it all the more difficult to get a firm grasp on the risk of these securities. This does not mean that investors should avoid them. They are priced fairly for their risk.

Collateralized Mortgage Obligations (CMOs) Another type of mortgage security is the collateralized mortgage obligation (CMO). CMOs, like mortgage pass-throughs, are debt securities whose payoffs are based on a portfolio of mortgages. CMOs, however, have risks and rewards that differ substantially from standard pass-throughs or IO and PO strips. In fact, some CMOs have considerably more risk and others have considerably less risk than pass-throughs or strips. This is accomplished by structuring the different types of CMO securities, called tranches, to have different priorities on the payments made on the underlying mortgages. In addition, some CMOs serve as equity, providing support to the liabilities and being paid after all other tranches are paid. This is called the residual tranche.

CMOs can be extremely complex securities, but we can understand the basics with a few simple examples. In the mortgage pass-through described above, the combined value of the IO and PO is \$120,944. Suppose we hold this mortgage and issue debt tranches with a total value of \$100,000, leaving a residual or equity tranche of \$20,944, which has the lowest priority of claims on the underlying mortgages. There are numerous variations of tranches in CMOs, distinguished by the different coupon rates and prepayment provisions. One type of tranche, called a planned amortization class or PAC, is characterized by a nearly fixed amortization schedule. In other words, it is designed to not be subject to prepayment risk. If, however, prepayments are dramatically above or below a specified reasonable range, some prepayments may be applied to the PAC tranche. Nonetheless, its substantially reduced prepayment risk results in a lower coupon.

With the PAC tranche protected against prepayment risk, obviously other tranches must bear more of that risk. The other tranches are arranged in priority according to which bears the greater prepayment risk. The tranches with higher coupons will receive larger percentages of any prepayments. In some cases, the tranche bearing the greatest prepayment risk must be fully paid off before the other tranches begin receiving

---

[4]The limited downside gain is often referred to as the negative convexity, a concept also associated with callable bonds.

prepayments. All tranches generally accrue interest, however, while higher-priority tranches are being paid off. In that manner the tranches accrue interest and are retired sequentially. Some tranches are called Z-bonds. They accrue interest at a given rate but no principal is paid until all other tranches are paid off.

All tranches tend to gain in value as interest rates decrease but the tranches bearing the greater prepayment risk gain less and the Z-bond gains more. It continues to accrue interest while the other tranches start receiving more prepayments. When interest rates increase, the securities bearing the greater prepayment risk do better because their prepayments are invested at higher rates. Though their values tend to fall, they fall less. The Z-bond value will then fall more.

When all tranches are paid off, there may be value left over, which is paid to the holders of the residual tranche. Any gain in value from the amount paid for the residual tranche will result from unexpected changes in interest rates and their effect on deviations in the prepayment rate from the assumptions built into the original structure.

Collateralized mortgage obligations are some of the most complicated and volatile financial instruments. Like mortgage pass-throughs, they have been aggressively marketed, sometimes to investors who had little understanding of the risks. Unfortunately, considerable misunderstanding occurred because many mortgage-backed securities were referred to as "government guaranteed." Of course, the guarantee is only that if a homeowner defaults, the payments will be made by the government agency, usually GNMA. Thus, this is only a credit guarantee; it does not guarantee against losses arising from interest rate changes and changes in the rate of prepayments. The volatility and illiquidity often seen in the market for mortgage-backed securities suggests that they be used only with caution and then only by investors who fully understand them.

The concept of securitizing mortgages has been extended to securitizing many other types of instruments, such as bonds and loans. In Chapter 15 we discuss credit derivatives, which can be combined with some of these other securitized debt instruments to create more opportunities for managing risk.

EXOTIC OPTIONS

In recent years, the proliferation of option products has led to a new class of options called exotic options. Although it is difficult to identify exactly where ordinary options end and exotic options begin, it is becoming common to refer to almost any option that is not traded on an exchange or not essentially identical to one traded on an exchange as an exotic option. In some cases these options are simple; in other cases, they are quite complex. What distinguishes them from what we have previously covered, however, is the fact that they offer different types of payoffs. What makes them like what we have previously covered is that their final payoffs are determined by whether a value exceeds or is less than an exercise price.

Digital and Chooser Options

Digital Options Digital options, which are sometimes called binary options, are actually very simple options and are of two types. An asset-or-nothing option pays the holder the asset or its equivalent cash value if the asset price at expiration exceeds the exercise price and nothing if the asset price at expiration is less than the exercise price. A cash-or-nothing option pays the holder a fixed amount of cash at expiration if the asset price at expiration exceeds the exercise price and nothing if the asset price at expiration is less than the exercise price. These descriptions apply to call versions, but put versions also exist. These options can be useful for hedging positions in standard European calls. In fact, there is a simple relationship between digital options and a standard European option.

Table 14.4 illustrates the payoffs to the holder of an asset-or-nothing option who also sells X cash-or-nothing options, each paying $1 if $S_T > X$. Notice that the payoffs are $S_T - X$ if $S_T > X$, and zero otherwise. Does this look familiar? It should because it is the payoff of a standard European call. This gives us some insights into how to price digital options. Recall that the Black-Scholes-Merton price is $S_0N(d_1) - Xe^{-r_cT}N(d_2)$. The first term is the price of the asset-or-nothing option. The second term, ignoring the minus sign, is the price of X cash-or-nothing options that pay off \$1 if they expire in-the-money. In other words, an asset-or-nothing option price is given as

$$O_{aon} = S_0N(d_1).$$

A standard cash-or-nothing option price is given as

$$O_{con} = e^{-r_cT}N(d_2),$$

so X cash-or-nothing options are worth $Xe^{-r_cT}N(d_2)$. As an example, consider an asset-or-nothing option written on the S&P 500 Total Return Index, which is at 1440. At expiration, the option pays the full cash value of the S&P 500. The exercise price is 1,440. The risk-free rate is 4.88 percent continuously compounded, the standard deviation is 11 percent, and the option expires in exactly one-half year. Since the option is written on the Total Return Index, which includes the dividends, we need not adjust for dividends when pricing the option. Substituting these values into d_1 for the Black-Scholes-Merton model gives $d_1 = 0.3526$. Using Table 5.1, we have N(0.35) = 0.6368. Then the asset-or-nothing option is worth 1,440(0.6368) = 917.00. If this option had been 1,440 cash-or-nothing options with an exercise price of 1,440, we would need d_2, which would be 0.2748. The N(0.27) = 0.6064 and the cash-or-nothing options would be worth, $1440e^{-0.0488(0.5)}(0.6064) = 852.17$. The difference in these two values, 64.83, is the value of a standard European call.

Table 14.4 Payoffs to a Position Long an Asset-or-Nothing Option and Short a Cash-or-Nothing Option

	Payoffs from Portfolio Given Asset Price at Expiration	
Payoff from	$S_T \le X$	$S_T > X$
Long asset-or-nothing option	0	S_T
Short X cash-or-nothing options	0	$-X$
Total	0	$S_T - X$

Recall that earlier in the chapter we looked at pay-later options, in which the buyer of the option pays the premium at expiration. This result is achieved by effectively borrowing the option premium. Thus, the option premium is paid, not at the start, but at expiration, where it is increased by the interest accrued over the life of the contract. A variation of this type of contract is the contingent-pay option. In this type of option, the premium is paid at expiration but only if the option expires in-the-money. To see how this option is identical to instruments we already know, consider the following combination. Buy an ordinary call expiring at T with exercise price X and sell C_{cp} cash-or-nothing calls, each paying \$1 at expiration. Of course we do not yet know the value of C_{cp}, which represents the premium of the contingent-pay option, but we can determine it. Table 14.5 shows the payoff of this combination at expiration, using a stock as the underlying.

TABLE 14.5 Payoffs from Contingent-Pay Option

		Payoffs from Portfolio Given Stock Price at Expiration	
Instrument	Current Value	$S_T \le X$	$S_T > X$
Long call	$C_e(S_0,T,X)$	0	$S_T - X$
Short C_{cp} cash-or-nothing calls	$-C_{cp}O_{con}$	0	$-C_{cp}$
	$C_e(S_0,T,X) - C_{cp}O_{con}$	0	$S_T - X - C_{cp}$

We see that this combination of a long call and C_{cp} cash-or-nothing calls each paying $1 at expiration replicates a contingent-pay option. Therefore, the value of the contingent pay option should equal the value of this combination at any time. We know, however, that at the start the value of a contingent-pay option is zero, because there is no payment up front. Recall that the value of a cash-or-nothing option paying $1 at expiration is $e^{-r_c T}N(d_2)$. Therefore, let us set the value of this combination at the start to zero:

$$C_e(S_0,T,X) - C_{cp}e^{-r_c T}N(d_2) = 0,$$

and solve for C_{cp}:

$$C_{cp} = \frac{C_e(S_0,T,X)}{e^{-r_c T}N(d_2)}.$$

This is the premium of the contingent-pay option that would be paid at expiration but only if the option expires in-the-money.

Let us work an example using the same digital option we looked at in the previous section. This option was based on the S&P 500 Total Return Index. Recall that the index is at 1,440 and the exercise price is also 1,440, the risk-free rate is 4.88 percent continuously compounded, the standard deviation is 11 percent, and the option expires in one-half year. We determined that the value of a standard European call is 64.83. The value of $N(d_2)$ is 0.6064. The premium of the contingent-pay option would be

$$C_{cp} = \frac{64.83}{e^{-(0.0488)(0.5)}0.6064} = 109.55.$$

Note that the contingent-pay premium is considerably higher than the premium of an ordinary European call. This is because the contingent-pay premium is deferred and could possibly never have to be paid at all.

Like a forward contract and a swap, the contingent-pay option has zero value at the start. During its life, however, its value changes and can be either positive or negative. Thus, we need a formula to determine its value. Suppose we are at a time t during the life of the option. Since the premium, C_{cp}, was set at the start, we know the payoff will be either $S_T - X - C_{cp}$ or zero. We can then value the contingent-pay option as a combination of an ordinary call and C_{cp} cash-or-nothing calls that pay $1 at expiration. In other words, we simply value it the same as the portfolio we constructed to solve for C_{cp}.

For example, suppose that two months into the life of this contingent-pay option, the index is at 1400. We using the following values as inputs: $S_t = 1{,}400$, $X = 1{,}440$, $r_c = 0.0488$, $T - t = 0.333$ (based on four months remaining), $\sigma = 0.11$, and $C_{cp} = 109.55$. Using a spreadsheet for greater accuracy, we find that $d_1 = -0.1561$, $N(d_1) = 0.4380$, $d_2 = -0.2195$, $N(d_2) = 0.4131$. The value of a standard European call would be 27.89. The value of 109.55 cash-or-nothing calls that each pay $1 at expiration is $109.55e^{-0.0488(0.333)}(0.4131) = 44.53$. Thus, the value of the contingent-pay call is $27.89 - 44.53 = -16.64$. As we noted, this type of option can have a negative payoff; hence, it can clearly have a negative value prior to its expiration.

Chooser Options A chooser option permits the investor to decide at a specific point during the life of the option whether it should be a call or a put. For example, suppose an investor believes the market will make a strong move but does not know which direction it will go. As we studied in Chapter 7, an investor can purchase a straddle, which is a call and a put. A cheaper alternative is to purchase a chooser option. The chooser allows the investor to decide at time t, which is prior to expiration, whether to make it a call or a put. Chooser options are also called as-you-like-it options.

Suppose at time t, an investor decides to make it a call. Then at expiration, the call pays off $S_T - X$ if $S_T > X$ and zero otherwise. If the investor had chosen the put, it pays off $X - S_T$ if $S_T \le X$ and zero otherwise. Note that unless $S_T = X$, a straddle will always pay off through either the call or the put, but it is possible that a chooser will not pay off at all.

Pricing a chooser option is quite simple. At time t, the investor will make it a call if the value of the call exceeds the value of the put. At t the call price can be expressed as $C(S_t, T - t, X)$ and the put price can be expressed as $P(S_t, T - t, X)$. A little algebra using put-call parity shows that the call is worth more and would be chosen if $S_t > X(1 + r)^{-(T-t)}$. The chooser can be replicated by simply holding an ordinary call expiring at T with exercise price of X and a put expiring at t with exercise price of $X(1 + r)^{-(T-t)}$. An end-of-chapter problem asks you to verify this statement.

Let us consider a chooser option on the DCRB stock. Recall that the stock price is 125.94, the exercise price is 125, the risk-free rate is 4.56 percent discrete (4.46 percent continuous), the time to expiration is 0.0959 (35 days), and the volatility is 83 percent. Let us assume the choice must be made in 20 days. First let us price an ordinary straddle. From Chapter 5, we used the Black-Scholes-Merton model and found the call to be worth 13.55 and the put to be worth 12.08. Therefore, the straddle would cost 25.63. The chooser would be worth the value of an ordinary call (13.55) and the value of a put expiring in 20 days ($t = 20/365 = 0.0548$ so $T - t = 0.0959 - 0.0548 = 0.0411$) with an exercise price of $125(1.0456)^{-0.0411} = 124.77$ and a volatility of 83 percent. Plugging into the Black-Scholes-Merton model gives a put value of 7.80. Thus, the chooser would cost $13.21 + 7.80 = 21.01$. The lower cost of the chooser over the straddle comes from the fact that there is a possibility that the payoff at expiration will be zero.

A chooser option that allows the holder to designate at *any* time before expiration whether it will be a call or a put is called a complex chooser. It is much more difficult to evaluate and we do not cover it here.

Path-Dependent Options

An important distinction among options is whether the path followed by the underlying asset price matters to the price of the option. For standard European options the path taken does not matter. They can be exercised only at expiration so it does not matter how the asset reached a certain price. Standard European options are said to be path-independent. For some options, however, the path taken does matter. American options are a classic case in point. Certain paths lead to early exercise. Though we did not refer to them this way in Chapters 4 and 5, we can certainly now call them path-dependent options.[5]

In recent years a family of new options, however, has become better known for its property of path-dependency. Some of these options are based on the average price or the maximum or minimum prices of the underlying asset during the option's life. Others have the property that they can expire worthless before expiration or can expire worthless even if they apparently are in-the-money at expiration. Today when one hears of path-dependent options, the reference is nearly always to these types of options, although these are certainly not all of the path-dependent types.

The binomial model provides a good framework within which to see how path-dependent options work. Table 14.6 presents a simple three-period binomial model where the current asset price is 50 and each period it can go up 25 percent or down 20 percent. The risk-free rate is 5 percent. We shall use an exercise price of 50. The table shows the binomial tree with the asset moving from 50 at time 0 to either 97.66, 62.50, 40, or 25.60 at time 3.

Also shown in the table is a listing of each of the eight paths the asset could take. Notice that it could go up-up-up, up-up-down, or any of six other possible paths. Recalling from our binomial model material in Chapter 4, the probability factor, p, is $(1 + r - d)/(u - d)$, which the table shows is 0.556. The probability of a given path is $p^{\#\text{ of times price goes up}}(1 - p)^{\#\text{ of times price goes down}}$. For example, the probability of going up, down, and then up is $0.552^2 0.444^1 = 0.137$. Also, keep in mind that this is not the true probability but only the risk

[5]The concept of path-independence versus path-dependence is illustrated by considering the following situation encountered in traveling. Suppose you are in Baltimore and need to fly to San Francisco. There are numerous ways to get from Baltimore to San Francisco, including direct flights and many possible connecting flights. If your objective is only to get from Baltimore to San Francisco, it does not matter what route you take. This would be consistent with path independence. If your objective is to get from Baltimore to San Francisco in the least amount of time and/or at the lowest cost, then certain routings, especially direct flights, are preferred. This would be consistent with path dependence.

Table 14.6 Path-Dependent Options Information

		Time 0	Time 1	Time 2	Time 3
S_0:	50	S_0	S_1	S_2	S_3
u:	1.25				97.66
d	0.80			78.13	
r:	0.05		62.50		62.50
X:	50	50.00		50.00	
p:	0.556[1]		40.00		40.00
1 − p	0.444			32.00	
					25.60

Path No.	Path	Probability[2]	S_0	S_1	S_2	S_3	Average Price	S_{max}	S_{min}	European Call Payoffs	European Put Payoffs
1	uuu	0.171	50.00	62.50	78.13	97.66	72.07	97.66	50.00	47.66	0.00
2	uud	0.137	50.00	62.50	78.13	62.50	63.28	78.13	50.00	12.50	0.00
3	udu	0.137	50.00	62.50	50.00	62.50	56.25	62.50	50.00	12.50	0.00
4	duu	0.137	50.00	40.00	50.00	62.50	50.63	62.50	40.00	12.50	0.00
5	udd	0.110	50.00	62.50	50.00	40.00	50.63	62.50	40.00	0.00	10.00
6	dud	0.110	50.00	40.00	50.00	40.00	45.00	50.00	40.00	0.00	10.00
7	ddu	0.110	50.00	40.00	32.00	40.00	40.50	50.00	32.00	0.00	10.00
8	ddd	0.088	50.00	40.00	32.00	25.60	36.90	50.00	25.60	0.00	24.40

European call price: 11.50
European put price: 4.69

[1]Computed as $(1 + r - d)/(u - d)$.
[2]Probability of a given path is $p^{\#\text{ of times price goes up}}(1 - p)^{\#\text{ of times price goes down}}$. For example, for path 5 (udd), the probability is $0.556^1 0.444^2 = 0.110$.

neutral probability, as discussed in Chapter 4. It is important not to combine paths, such as the fact that the stock can reach 62.50 at expiration three different ways. Each path can lead to different values of path-dependent options, so we must keep the paths separate.

The table also shows the sequence of asset prices for each path: S_0, S_1, S_2, and S_3. The eighth column shows the average price over the path, counting the original asset price. Thus, for path 1, the average is (50.00 + 62.50 + 78.13 + 97.66)/4 = 72.07. The next two columns show the maximum and minimum prices attained over the given path.

The last two columns show the payoffs at expiration for standard European calls. The value for a given path is $\text{Max}(0,S_3 - 50)$ for calls and $\text{Max}(0, 50 - S_3)$ for puts, which reflects those options' dependencies on only the final asset price, S_3. To find the price of those options, we simply multiply each option payoff by the probability of that payoff occurring, sum the payoffs, and divide by $(1.05)^3$, which is simply discounting the total to the present. This is the same procedure we used in previous applications of the binomial model in Chapter 4. The European call price is 11.50 and the European put price is 4.69.[6] We now look at several types of path-dependent options.

Asian Options An Asian option is an option where the payoff is based on the average price attained by the underlying asset during the life of the option or over a specified period of time, which does not have to correspond to the life of the option. An Asian option can take the form of either a call or a put. The average price can be used in place of either the expiration price of the asset or the exercise price. The former are called average price options while the latter are called average strike options. For the three-period case, the payoffs at expiration are as follows:

[6]As an example, the call price is found as $(0.171(47.66) + 0.137(12.50) + \cdots + 0.088(0))(1.05)^3 = 11.50$.

Average price call: $Max(0, S_{avg} - X)$

Average put price: $Max(0, X - S_{avg})$

Average strike call: $Max(0, S_3 - S_{avg})$

Average strike put: $Max(0, S_{avg} - S_3)$.

Table 14.7 illustrates how to price an Asian option. Look at the fifth column, which shows the payoffs for the average price call. For path 1, at expiration the average price is 72.07, which makes the option be in-the-money by 22.07. The remaining payoff values are computed similarly using the first formula above. Then the values in the payoff column are multiplied by their respective probabilities (column 3). The sum of these probability-weighted outcomes is then discounted over three periods to the present. We see that both types of Asian options are considerably less valuable than their European counterparts. There are several reasons for this, but the most obvious one is that the volatility of an average of a series of numbers is less than the volatility of the series itself.

Table 14.7 Pricing an Asian Option

Path No.	Path	Probability	Average Price (S_{avg})	Average Price Call Payoff[1]	Average Price Put Payoff[2]	Average Strike Call Payoff[3]	Average Strike Put Payoff[4]
1	uuu	0.171	72.07	22.07	0.00	25.59	0.00
2	uud	0.137	63.28	13.28	0.00	0.00	0.78
3	udu	0.137	56.25	6.25	0.00	6.25	0.00
4	duu	0.137	50.63	0.63	0.00	11.88	0.00
5	udd	0.110	50.63	0.63	0.00	0.00	10.63
6	dud	0.110	45.00	0.00	5.00	0.00	5.00
7	ddu	0.110	40.50	0.00	9.50	0.00	0.50
8	ddd	0.088	36.90	0.00	13.10	0.00	11.30
			Value of Option[5]	5.72	2.37	5.94	2.48

[1] $Max(0, S_{avg} - 50)$.
[2] $Max(0, 50 - S_{avg})$.
[3] $Max(0, S_3 - S_{avg})$.
[4] $Max(0, S_{avg} - S_3)$.
[5] Value of Option is found as the sum of the product of the payoffs of the option (column 5, 6, 7, or 8) by the respective probability of that payoff (column 3) discounted at the risk-free rate (i.e., divided by 1.05^3).

The lower prices of Asian options relative to European options make them useful for situations in which the user wants to hedge or speculate on the price series but is unwilling to spend the full price required by European options. In some cases a hedge of only an average price is satisfactory. Asian options are also used in markets where prices are unusually volatile and may be susceptible to distortion or manipulation. The averaging process results in the option payoff not being completely dependent on the price at expiration.

As we discussed in Chapter 4, the binomial model requires a large number of time steps to give an accurate approximation to the Black-Scholes-Merton price for standard European options. Unfortunately there is no formula like the Black-Scholes-Merton for Asian options.[7] The binomial model is useful for illustrative purposes, but applying it in practice with a large number of time steps is impractical for Asian options. The computer must keep track of 2^n paths, where n is the number of time periods. For our three-

[7] An exception is that if the average is a geometric average, then a formula does exist. Geometric average Asian options, however, are rarely used in practice.

period model, the number of paths is $2^3 = 8$. For a realistic situation, we would require at least 50 time steps, in which case the computer must keep track of 1,125,900,000,000,000 paths! There are some Asian option approximation formulas used in practice, but a convenient method of pricing Asian options is by Monte Carlo simulation. We illustrate the Monte Carlo simulation procedure in Appendix 14.

Lookback Options Another interesting path-dependent option is the lookback option. A lookback call allows you to buy the asset at the lowest price and a lookback put allows you to sell it at the highest price attained during the life of the option. It is said that lookback options allow you to "buy low, sell high." Since you never have to miss in timing the market, a lookback option is also called a no-regrets option.

For the three-period binomial example, the standard lookback option has payoffs at expiration as follows:

$$\text{Lookback call: Max}(0, S_3 - S_{min})$$

$$\text{Lookback put: Max}(0, S_{max} - S_3).$$

In other words, the lookback call sets the exercise price at the lowest price the asset attained during the option's life. A lookback put sets the exercise price at the highest price it attained during the life of the option. In the worst case, a lookback option expires at-the-money.

Another class of lookback options, which are sometimes called modified lookback options, has the exercise price fixed but replaces the terminal asset price with the maximum price for a call and minimum price for a put. Thus, their payoffs are:

$$\text{Fixed strike lookback call: Max}(0, S_{max} - X)$$

$$\text{First strike lookback put: Max}(0, X - S_{min}).$$

Table 14.8 illustrates the pricing of lookback options. Consider the standard lookback call, whose payoffs are listed in column 6. If path 1 occurs, the asset ends up at 97.66. The minimum price along that path is 50. Thus, it expires in-the-money worth 47.66. Notice particularly path 7. From Table 14.6, we see that the asset price ended up at 40 but it went as low as 32, so the call ends up worth 8. Similarly, the basic lookback put in path 2 pays off 15.63 because the asset ended up at 62.50 but went as high as 78.13. The remaining payoffs are determined in a similar manner using the appropriate formulas shown above. The value of each

Table 14.8 Pricing a Lookback Option

Path No.	Path	Probability	Maximum Price (S_{max})	Minimum Price (S_{min})	Lookback Call Payoff[1]	Lookback Put Payoff[2]	Fixed Strike Lookback Call Payoff[3]	Fixed Strike Lookback Put Payoff[4]
1	uuu	0.171	97.66	50.00	47.66	0.00	47.66	0.00
2	uud	0.137	78.13	50.00	12.50	15.63	28.13	0.00
3	udu	0.137	62.50	50.00	12.50	0.00	12.50	0.00
4	duu	0.137	62.50	40.00	22.50	0.00	12.50	10.00
5	udd	0.110	62.50	40.00	0.00	22.50	12.50	10.00
6	dud	0.110	50.00	40.00	0.00	10.00	0.00	10.00
7	ddu	0.110	50.00	32.00	8.00	10.00	0.00	18.00
8	ddd	0.088	50.00	25.60	0.00	24.40	0.00	24.40
			Value of Option[5]	13.45	7.73	14.54	6.64	

[1] $\text{Max}(0, S_3 - S_{min})$.
[2] $\text{Max}(0, S_{max} - S_3)$.
[3] $\text{Max}(0, S_{max} - 50)$.
[4] $\text{Max}(0, 50 - S_{min})$.
[5] Value of Option is found as the sum of the product of the payoffs of the option (column 6, 7, 8, or 9) multiplied by the respective probability of that payoff (column 3) discounted at the risk-free rate (i.e., divided by 1.05^3).

option is once again found by multiplying each option payoff by the appropriate probability (column 3), summing these probability-weighted payoffs, and discounting at 5 percent over three periods.

Because lookback options allow their holders to achieve higher payoffs for some outcomes and positive payoffs for some outcomes where European options would have zero payoffs, they are worth substantially more. Is the additional price worth it? That depends on the investor. An investor who experiences anguish over failing to buy when the price is low or sell when the price is high would probably find the extra cost worth it.

Formulas for the prices of most lookback options have been derived. Using these formulas in practice, however, is complicated by the fact that many lookback options do not provide for continuous monitoring of the asset price. Instead they consider only the closing price of the asset in determining the maximum or minimum. Because of the aforementioned large number of paths that must be followed in the binomial model, Monte Carlo methods are often used.

Barrier Options The final type of path-dependent option we examine is the barrier option. While lookback options have positive payoffs for more outcomes than do standard European options, barrier options have fewer positive payoffs than European options. This makes them less expensive than standard European options. Investors who want to save a little money and do not want to pay for outcomes they do not believe are likely to occur will be attracted to barrier options.

In one type of barrier option the holder specifies a level of the asset price that, if touched by the asset before the option expires, causes the option to terminate. In the other type of barrier option, the holder specifies a level of the asset price that must be hit by the asset to activate the option. If the asset never hits the barrier during the option's life, the option cannot be exercised even if it expires in-the-money at expiration. Options that can terminate early are called out-options and sometimes knock-out options. Options that will be unexercisable if they fail to touch the barrier are called in-options or sometimes knock-in options. Out-options, thus, can expire prematurely with no value if the barrier is hit. In-options must be activated by hitting the barrier. If they are never activated, they automatically expire worthless. If the barrier is set above the current asset price, both in- and out-options are referred to as up-options. If the barrier is set below the current asset price, in- and out-options are referred to as down-options.

Table 14.9 illustrates the pricing of barrier options. The table shows the payoffs at expiration of all eight types of barrier options. We first look at down-and-out and down-and-in options. The bold zero values are

Table 14.9 Pricing a Barrier Option

			Down-and-Out and Down-and-In Options (H = barrier)				Up-and-Out and Up-and-In Options (H = barrier)			
Path No.	Path	Probability	Out Call Payoff (H = 45)	In Call Payoff (H = 45)	Out Put Payoff (H = 35)	In Put Payoff (H = 35)	Out Call Payoff (H = 70)	In Call Payoff (H = 70)	Out Put Payoff (H = 55)	In Put Payoff (H = 55)
1	uuu	0.171	47.66	0.00	0.00	0.00	0.00	47.66	0.00	0.00
2	uud	0.137	12.50	0.00	0.00	0.00	0.00	12.50	0.00	0.00
3	udu	0.137	12.50	0.00	0.00	0.00	12.50	0.00	0.00	0.00
4	duu	0.137	0.00	12.50	0.00	0.00	12.50	0.00	0.00	0.00
5	udd	0.110	0.00	0.00	10.00	0.00	0.00	0.00	0.00	10.00
6	dud	0.110	0.00	0.00	10.00	0.00	0.00	0.00	10.00	0.00
7	ddu	0.110	0.00	0.00	0.00	10.00	0.00	0.00	10.00	0.00
8	ddd	0.088	0.00	0.00	0.00	24.40	0.00	0.00	24.40	0.00
	Value of Option[1]		10.02	1.48	1.90	2.80	2.96	8.54	3.75	0.95

[1]Value of Option is found as the sum of the product of the payoffs of the option (columns 4, 5, 6, 7, 8, 9, 10, or 11) by the respective probability of that payoff (column 3) discounted at the risk-free rate (i.e., divided by 1.05^3).

Note: Bold values for out options are outcomes in which the option hit the barrier and, thus, was knocked out. Italicized values for in-options are outcomes in which the option failed to hit the barrier and, thus, was not exercisable at expiration.

outcomes in which out-options hit the barrier and, thus, were knocked out with no value. The italicized values are outcomes in which in-options failed to hit the barrier and, thus, were never activated.

Consider the down-and-out call illustrated in column 4 where the barrier is set at 45. This means that if the asset ever hits 45, the option terminates with no value. This will occur in paths 4, 6, 7, and 8 where in each path the asset hit 40 at time 1. In path 5 the option hits the barrier but expires worthless anyway. Now move one column over to the down-and-in call. This option will never activate if the asset does not fall to 45 during the option's life. This occurs in paths 1, 2, and 3.

The values of the options are shown in the bottom row of the table and are computed the same way as in the previous examples. Notice a very interesting result. If one combines the payoffs from the down-and-out call with those of the down-and-in call, one obtains the same payoffs of a standard European call. This should make sense. When one holds both options, if the out-option knocks out, the in-option immediately knocks in. Thus, the value of the two options combined is $10.02 + 1.48 = 11.50$, which is the price we previously obtained for the standard European call. Likewise the down-and-out put and down-and-in put add up in value to the value of the standard European put, subject to a round-off error. Similar results hold for the up-options.

Some barrier options have an additional feature called a rebate. If a knock-out barrier option provides for a rebate, the holder is paid a fixed sum of cash whenever the option knocks out. If a knock-in option provides for a rebate, the holder is paid a fixed sum of cash at expiration if the option never crosses the barrier. When a barrier option has a rebate, the values of the complementary out- and in-options will add up to more than the standard European option price.

As noted above, barrier options allow investors to avoid paying a price for states they do not believe will occur. As an example of the use of a barrier option, consider a protective put. A standard European put would cost 4.69. To save a little money, an investor might purchase an up-and-out protective put with a barrier of 55 for 3.75. This option would knock out if the asset goes to 55. The holder would in effect be saying that if the asset price rises quickly, he does not expect to need the insurance. Thus, his willingness to terminate the insurance early allows him to save money. Note, however, that this does not guarantee a better result. It simply eliminates the payoff in path 5. The knocked-out put will not pay off if the asset goes up, down, and then down because it knocked out at time 1, even though at expiration the asset price is below the exercise price. Similarly an investor interested in holding a call might consider purchasing a down-and-out call. The standard call will cost 11.50 while the down-and-out with the barrier at 45 will cost 10.02. The barrier option will not pay off in path 4 where the asset goes down, up, and up whereas the standard call would pay off if that occurred.

There are many variations of path-dependent options, including some that are American-style, that is, with early exercise. Barrier options in particular can have many types of twists, including barriers both above and below the current price and requirements that the barrier be hit more than once. Another type of barrier option allows the holder to reset the exercise price if the barrier is hit. A more common and somewhat complicating factor is that often only the closing price is considered when determining whether the barrier was touched. This necessitates the use of the binomial model with a large number of time steps and special care to position the barrier on the nodes. Monte Carlo methods, as described in Appendix 14, are also required.

Other Exotic Options

Most of the exotic options described above are reasonably simple to understand. There are many others that are quite complex and we will not attempt to tackle the pricing of these instruments. It is useful to examine the basic characteristics of these instruments so that you will be aware of them and perhaps recognize them when you encounter them in the future.

Compound options are options on options. For example, you can buy a call option that gives you the right to buy another call option. This is a call on a call. Likewise there are calls on puts, puts on calls, and puts on puts. The latter two give you the right to sell a call or a put. Compound options would be useful if an investor thought he might need an option later and wanted to establish a price at which the option could be bought or sold. A variation of the compound option is the installment option, which permits the premium to be paid in equal installments over the life of the option. At each payment date, the holder of the option decides whether it is worth paying the next installment. This is equivalent to deciding whether it is worth exercising the compound option to acquire another option.

Multi-asset options consist of a family of options whose payoffs depend on the prices of more than one asset. A simple type of multi-asset option is the exchange option. This option allows the holder to acquire one asset by giving up another. This kind of option is actually a special case of an ordinary European call. In a European call, the call holder has the right to give up one asset, cash, and acquire another asset, the stock. Exchange options are not commonly traded in the over-the-counter market, but they have proven to be useful in valuing many other types of options, such as the quality option in the Treasury bond futures contract that we covered in Chapter 10.

Another type of multi-asset option is the min-max option. This is an option that pays off according to the better or worse performing of two assets. For example, consider a call on two assets, stocks X and Y, in which at expiration, we determine which asset has the higher price and then calculate the payoff based on that asset. This is an option on the maximum of two assets. In an option on the minimum of two assets, the payoff is calculated based on the asset with the lower price at expiration. These options can be puts or calls, and they are sometimes called rainbow options. A slight and more common variation determines the better or worse performing asset, not on the basis of how high or low its price is, but according to the rate of return on the asset during the option's life. This is also called an alternative option. Another variation of this is the outperformance option, whose payoff is determined by the difference in the prices or rates of return on two assets or indices relative to an exercise price or rate.

Shout options permit the holder at any time during the life of the option to establish a minimum payoff that will occur at expiration. For example, suppose the stock price is very high relative to the exercise price. The investor, holding such an option and fearing that the stock price will fall, can establish the minimum payoff as the amount by which it is currently in-the-money. A slight variation of this is a cliquet option, in which the exercise price can periodically increase as the stock price rises. Still another variation, the lock-in option, permits establishment of the exact payoff, as opposed to just the minimum, prior to expiration. A deferred strike option is one in which the exercise price is not set until a specific date prior to expiration.

Forward-start options are options whose lives do not begin until a later date. The premium is paid up front and the purchaser specifies a desired degree of moneyness, such as at-the-money, 5 percent out-of-the-money, and so forth. Once the option begins, it is like an ordinary option. This type of option is similar to the types of executive and employee stock option plans used by firms. The firm makes a commitment that will result in its awarding options at various future dates. When the options are structured so that as one expires another begins, the combination is called a tandem option.

The exotic options described here are the primary ones that are used in today's markets, but there are many more. A few years ago, most of these options did not exist. The pace of creativity is quite rapid. As your career in the financial world evolves, you will likely encounter new ones quite often.

SOME UNUSUAL DERIVATIVES

In most of this book we have examined derivatives based only on stocks, bonds, and currencies. On occasion we have made reference to derivatives on certain commodities such as metals and oil. In recent years, derivatives have emerged on some most unusual underlyings. We take a look at two of these here.

Electricity Derivatives

Most of the energy consumed on the earth derives from three primary sources: fossil fuels, nuclear reaction, and the sun. Fossil fuels in the form of oil, coal, and natural gas are the primary sources of energy. We have already mentioned derivatives on oil, and derivatives on natural gas exist as well. Derivatives on coal have not, however, developed to any great extent. These sources of energy are also converted into other sources of energy, the primary one of which is electricity, which is obviously quite widely used in businesses and by consumers. Until recent years, electricity was largely regulated by the various states in the U.S. and foreign countries, but these governments have begun to relax this regulation, thereby allowing electricity prices to fluctuate and giving consumers the choice of where to buy their electricity. In response to this new-found price volatility, the energy industry created spot and derivative markets on electricity. These markets have proven to be useful by large-volume electricity users and public utilities for managing the volatility of electricity prices.

On the surface it would seem that one could use standard models for pricing forwards, futures, and options on electricity, but that is not the case. Electricity is quite unlike any asset we have studied so far in that it cannot be stored. Stocks, bonds, currencies, and most commodities can be purchased and held, but electricity is manufactured, sent along lines to where it is needed, and almost immediately consumed. All of the models we have learned are based on the idea that the underlying asset can be held for a finite period of time. Hence, we have reason to believe that the cost of carry model and the Black-Scholes-Merton model cannot be used to price electricity derivatives. Indeed the industry generally does not use these models. Unfortunately there is little agreement on a better alternative. These issues are quite advanced for this level of the subject, and we are not covering them here. In spite of the lack of agreement on pricing electricity derivatives, the market for electricity derivatives continues to grow. Confusion over pricing clearly does not stop a market in which the product serves an important need. In fact, confusion over pricing could even stimulate trading if some investors believe they know the proper way to price these instruments before everyone else knows.

Weather Derivatives

All of the underlyings we have covered can be owned. Even electricity can be owned, even though as we noted, it is consumed immediately. But in general, the most important feature of an underlying is not that it can be owned, but that it is a source of risk. One important source of risk is the weather. There are few companies whose operations are completely unaffected by the weather. Some entities' fortunes are tied almost completely to the weather. For example, the profitability and survival of orange juice growers in central Florida are highly affected by temperatures, primarily the potential for freezes. Farmers in general are highly exposed to weather. But so are airlines, ski resorts, various tourist attractions, and recreational services. Retail establishments and public utilities are likewise affected by the weather.

To manage these risks an array of options, forwards, futures, and swaps on the weather have been created in recent years. Weather has many excellent properties that make it appropriate for having derivatives created thereon. One important justification for weather derivatives is that weather is highly measurable, and there is a long history of data on the weather. There are three primary ways to measure the weather. One is by temperature, one by precipitation, and the third by the financial damage caused by the weather.

Information on temperature is widely available. The weather derivatives industry has created a standard measure used in derivative contracts called the heating degree day and cooling degree day. In the U.S., a benchmark of 65 degrees Fahrenheit is considered to be a normal level of comfort. On a cool (<65 degrees) day, energy is considered to be consumed in providing heat. A quantity of one heating degree day (HDD) is the difference between 65 degrees and the average temperature of that day. On a warm (>65 degrees) day, energy is considered to be consumed in providing air conditioning. A quantity of one cooling degree day (CDD) is the difference between the average temperature of that day and 65 degrees. Over a period of time, such as the life of a derivative contract, the number of heating degree days or cooling degree days is accumulated and compared to the exercise price of an option, the price of a forward or futures contract, or the fixed price of a swap.

Another measure of the weather is the quantity of precipitation. Suppose, for example, that a ski resort knows that unless it receives at least 10 feet of snow over the December through March period, it will not achieve its target cash flow. It might buy a put option on the amount of snowfall. The option pays off based on the difference between 10 feet and the actual amount of snow received. If the quantity of snow received exceeds 10 feet, the option would expire out-of-the-money. Of course, the ski resort could sell a forward contract, which would not require a payment up front, but would require the resort to pay off based on the difference between the actual amount of snowfall and the benchmark amount, which would not necessarily be 10 feet.

A third measure of weather is the amount of financial loss incurred from weather damage. Hurricanes, earthquakes, tornadoes, and floods are the primary forms of weather damage. The insurance industry tabulates losses based on its claims. Derivative contracts are based on these financial loss figures. The insurance industry itself is particularly attracted to these types of derivatives, because it gives the industry a means of shedding some of its risk.

Pricing weather derivatives, like pricing electricity derivatives, is particularly challenging. Weather cannot be stored, but storable assets are influenced by the weather and the effects of weather can be measured. Suppose that the aforementioned ski resort can predict with relative certainty the number of skiers it will have for a various amounts of snow over the season. It might determine that if it receives only 10 feet of snow, it will generate sufficient cash to earn a risk-free return. Thus, a forward contract priced at 10 would earn a risk-free return. This would form the basis for pricing the contract. A similar, though slightly more complex analysis, would be used to price an option. Of course, not all weather derivatives can be priced easily, but as noted, weather data are available in large quantities over long periods of time, which facilitates measuring its effects on the cash flows of a business.

The weather derivatives industry has grown slowly but steadily. Interestingly, it has offered many new job opportunities to meteorologists. No longer are they primarily limited to employment at local television stations. Many of these weather experts now work for financial institutions and commodity trading companies.

QUESTIONS AND PROBLEMS

1. Explain the advantages and disadvantages of implementing portfolio insurance using stock and puts in comparison to using stock and futures in a dynamic hedge strategy.
2. Explain how a portfolio manager might justify the purchase of an inverse floating-rate note.
3. Demonstrate that the payoffs of a chooser option with an exercise price of X and a time to expiration of T that permits the user to designate it as a call or a put at t, can be replicated with two transactions. Specifically, by (1) buying a call with an exercise price of X and time to expiration of T and (2) buying a put with an exercise price equal to $X(1 + r)^{-(T-t)}$ and time to expiration of t. This proof will require that you consider two possible outcomes at t (user designates it as a call or user designates it as a put according to the rule given in this chapter). For each outcome at t, there are two possible outcomes at T, $S_T \geq X$ or $S_T < X$. Explain why a chooser option is less expensive than a straddle.
4. Explain why an interest-only (IO) mortgage strip has a value that is extremely volatile with respect to interest rates. What two factors determine its value?
5. Explain the difference between path-dependent options and path-independent options and give examples of each.
6. Give an example of a situation in which someone might wish to use a barrier option.
7. Explain how weather derivatives could be used by an electric utility to manage the risk associated with power consumption as affected by the weather.
8. In modern financial derivatives markets, there are many exotic options. Briefly explain compound options, multi-asset options, shout options, and forward start options.
9. On July 5 a market index is at 492.54. You hold a portfolio that duplicates the index and is worth 20,500 times the index. You wish to insure the portfolio at a particular value over the period until September 20. You can buy risk-free debt maturing on September 20 with a face value of $100 for $98.78.

a. You plan to use puts, which are selling for \$23.72 and have an exercise price of 510. Determine the appropriate number of puts and shares to hold. What is the insured value of the portfolio?
b. Determine the value of the portfolio if the index on September 20 is at 507.35.
c. Determine the value of the portfolio if the index on September 20 is at 515.75. Compute the upside capture and the cost of the insurance.

10. Use the information in problem 9 to set up a dynamic hedge using stock index futures. Assume a multiplier of 500. The futures price is 496.29. The volatility is 17.5 percent. The continuously compounded risk-free rate is 3.6 percent, and the call delta is 0.3826. Let the stock price increase by \$1, and show that the change in the portfolio value is almost the same as it would have been had a put been used.

For the next three problems, use a two-period binomial model on a stock worth 100 that can go up 20 percent or down 15 percent. The risk-free rate is 6 percent each period.

11. Determine the price of an average price Asian call option. Use an exercise price of 95. Count the current price in determining the average. Comment on whether you would expect a standard European call to have a lower or higher price.
12. Determine the prices of lookback and modified lookback calls and puts. For the modified lookbacks, use an exercise price of 95.
13. Determine the prices of the following barrier options.
    a. A down-and-out call with the barrier at 90 and the exercise price at 95
    b. An up-and-out put with the barrier at 110 and the exercise price at 105
    c. Select any other barrier option but base your selection on the following instructions: Calculate the value of your selected barrier option and use it with the results you obtained in part a or b to determine the price of a standard European call or put. Then calculate the actual value of the European call or put and compare that answer with your answer obtained from the barrier options. Explain why this result is obtained.
14. A portfolio manager is interested in purchasing an instrument with a call option-like payoff but does not want to have to pay money up front. The manager learns from a banker that one can do this by entering into a break forward contract. The manager wants to learn if the banker is quoting a fair price. The stock price is 437.55. The contract expires in 270 days. The volatility is 18 percent and the continuously compounded risk-free rate is 3.75 percent. The exercise price will be set at the forward price of the stock.
    a. Determine the exercise price.
    b. The loan implicit in the break forward contract will have a face value of 40.19. Determine if this is a fair amount by using your answer in a and computing the value of K.
    c. Regardless of whether the break forward is found to be fairly priced, determine the value of the position if the stock price ends up at 465 and at 425.
15. Consider a stock priced at 100 with a volatility of 25 percent. The continuously compounded risk-free rate is 5 percent. Answer the following questions about various options, all of which have an original maturity of one year.
    a. Find the premium on an at-the-money pay-later call option. Then determine the market value of the option nine months later if the stock is at 110.
    b. Find the value of F and K on a break forward contract. Then determine the market value of the break forward nine months later if the stock is at 110.
    c. Find the premium on an at-the-money contingent-pay call option. Then determine the market value of the option nine months later if the stock is at 110.
16. A stock is priced at 125.37, the continuously compounded risk-free rate is 4.4 percent, and the volatility is 21 percent. There are no dividends. Answer the following questions.

a. Determine a fair price for a two-year asset-or-nothing option with exercise price of 120.
b. Assuming you purchased the asset-or-nothing option at the price you determined in *a*, calculate your profit if the asset price at expiration is (1) 138 and (2) 114.
c. Determine a fair price for a two-year cash-or-nothing option with exercise price of 120 that pays 120 if it expires in-the-money.
d. Assuming you purchased the cash-or-nothing option at the price you determined in c, calculate your profit if the asset price at expiration is (1) 138 or (2) 114.

17. Consider a 10-year, fixed-rate mortgage of $500,000 that has an interest rate of 12 percent. For simplification assume that payments are made annually.
a. Determine the amortization schedule.
b. Using your answer in a, determine the value of both IO and PO strips with a discount rate of 10 percent under the assumption that the mortgage will not be prepaid.
c. Now recompute the values of the IO and PO under the assumption that interest rates immediately fall to 8 percent and the mortgage is prepaid in year 6.
d. Explain the risk characteristics of IO and PO strips.

18. An investment manager expects a stock to be quite volatile and is considering the purchase of either a straddle or a chooser option. The stock is priced at 44, the exercise price is 40, the continuously compounded risk-free rate is 5.2 percent, and the volatility is 51 percent. The options expire in 194 days. The chooser option must be declared a call or a put exactly 90 days before expiration.
a. Determine the prices of the straddle and the chooser.
b. Suppose at 90 days before expiration, the stock is at 28. Find the value of the chooser option at expiration if the stock price ends up at 50 and at 30.
c. Suppose at 90 days before expiration, the stock is at 60. Find the value of the chooser option at expiration if the stock price ends up at 50 and at 30.
d. Compare your answers in *c* and *d* to the performance of the straddle.

19. Suppose FRM, Inc. issued a zero-coupon, equity index-linked note with a five-year maturity. The par value is $1,000 and the coupon payment is stated as 75% of the equity index return or as zero. Calculate the cash flow at maturity assuming the equity index appreciates by 30% over this five-year period.

20. (Concept Problem) Suppose you are asked to assist in the design of an equity-linked security. The instrument is a five-year zero coupon bond with a guaranteed return of 1 percent, compounded annually. At the end of five years the bond will pay an additional return based on any appreciation of the Nikkei 300 stock index, a measure of the performance of 300 Japanese stocks. The risk-free rate is 5.5 percent, compounded annually, and the volatility of the index is 15 percent. In addition the index pays a dividend of 1.7 percent continuously compounded. Presently the index is at 315.55 and the additional return is based on appreciation above the current level of the index. You expect to sell these bonds in minimum increments of $100. Overall you expect to sell $10 million of these securities. Your firm has determined that it needs a margin of $175,000 in cash today to cover costs and earn a reasonable profit. Determine the percentage of the Nikkei return that your firm should offer to cover its costs. Your firm would then set the percentage offered at less than this. If your firm sells this security, comment on the risk it creates for itself and suggest how it might deal with that risk.

21. (Concept Problem) A convertible bond is a bond that permits the holder to turn in the bond and convert it into a certain number of shares of stock. Conversion would, thus, occur only when the stock does well. As a result of the option to convert the bond to stock, the coupon rate on the bond is lower than it otherwise would be. A new type of financial instrument, the reverse convertible, pays a higher-than-normal coupon, but the principal payoff can be reduced if the stock falls. Let us specify that the principal payoff of the reverse convertible is FV, the face value, if $S_T > S_0$ where S_0 is the stock price when the bond is issued. If $S_T \leq S_0$, the principal payoff is $FV(S_T/S_0)$. Thus, for example, if the stock falls by 10 percent, S_T/S_0, the principal payoff, is 0.9FV. Show that this payoff (FV if $S_T > S_0$, and $FV(S_T/S_0)$ if S_T S_0) is equivalent to a combination of an ordinary bond and a certain number of European puts with an exercise price of S_0. Determine how many puts you would need.

Monte Carlo Simulation

Simulation is a procedure in which random numbers are generated according to probabilities assumed to be associated with a source of uncertainty, such as a new product's sales or, more appropriately for our purposes, stock prices, interest rates, exchange rates, or commodity prices. Outcomes associated with these random drawings are then analyzed to determine the likely results and the associated risk. Often this technique is called *Monte Carlo simulation,* being named for the city of Monte Carlo, which is noted for its casinos.

The gambling analogy notwithstanding, Monte Carlo simulation is a legitimate and widely used technique for dealing with uncertainty in many aspects of business operations. For our purposes, it has been shown to be an accurate method of pricing options and particularly useful for path-dependent options and others for which no known formula exists.

To facilitate an understanding of the technique, we shall look at how Monte Carlo simulation has been used to price standard European options. Of course, we know that the Black-Scholes-Merton model is the correct method of pricing these options, so Monte Carlo simulation is not really needed. It is useful, however, to conduct this experiment because it demonstrates the accuracy of the technique for a simple option of which the exact price is easily obtained from a known formula.

The assumptions of the Black-Scholes-Merton model imply that for a given stock price at time t, simulated changes in the stock price at a future time can be generated by the following formula:

$$\Delta S = Sr_c\Delta t + S\sigma\varepsilon\sqrt{\Delta t},$$

where S is the current stock price, S is the change in the stock price, r_c is the continuously compounded risk-free rate, is the volatility of the stock, and Δt is the length of the time interval over which the stock price change occurs. The variable is a random number generated from a standard normal probability distribution. Remember from Chapter 5 that the standard normal random variable has a mean of zero and a standard deviation of 1.0, and occurs with a frequency corresponding to that associated with the famous bell-shaped curve.

Generating future stock prices according to the above formula is actually quite easy. A standard normal random variable can be approximated with a slight adjustment to Microsoft Excel's Rand() function. The Rand() function produces a uniform random number between 0 and 1, meaning that it generates numbers between 0 and 1 with equal probability. A good approximation for a standard normal variable is obtained by the Excel formula "= Rand() + Rand() + Rand() + Rand() + Rand() + Rand() + Rand() + Rand() + Rand() + Rand() + Rand() + Rand() − 6.0," or simply 12 uniform random numbers minus 6.0.* Alternatively, a normal random number generator is available in Excel's Data Analysis tool pack.

---

*This approximation is based on the fact that the distribution of the sum of 12 uniformly distributed random numbers between 0 and 1 will have a mean of six and a standard deviation of 1. By subtracting 6.0, we adjust the mean to zero without changing the standard deviation. What we obtain is technically not normally distributed but is symmetric with a mean of zero and a standard deviation of 1.0, which are three properties associated with the normal distribution. The procedure is widely accepted as a quick and reasonable approximation but would not pass the most demanding tests for normality.

After generating one standard normal random variable, you then simply insert it into the right-hand side of the above formula for ΔS. This gives the price change over the life of the option, which is then added to the current price to obtain the price of the asset at expiration. You then compute the price of the option at expiration according to the standard formulas, $Max(0, S_T - X)$ for a call or $Max(0, X - S_T)$ for a put. This produces one possible option value at expiration. You then repeat this procedure many thousands of times, take the average value of the option at expiration, and discount that value at the risk-free rate. Some users compute the standard deviation of the call prices in order to obtain a feel for the possible error in estimating the price.

Let us price the DCRB June 125 call that we first encountered in Chapter 3. The DCRB stock price is 125.94, the exercise price is 125, the risk-free rate is 4.46 percent, the volatility is 83 percent, and the time to expiration is 0.0959 years. Inserting the above approximation formula for a standard normal random variable in any cell in an Excel spreadsheet produces a random number. Suppose that number is 0.733449. Inserting into the formula for ΔS gives $125.94(0.0446)(0.0959) + 125.94(0.83)(0.733449) = 24.28$. Thus, the simulated value of the stock at expiration is $125.94 + 24.28 = 150.22$. At that price, the option will be worth $Max(0, 150.22 - 125) = 25.22$ at expiration. We then draw another random number. Suppose we get -0.18985. Inserting into the formula for ΔS, we obtain $125.94\ (0.0446)(0.0959) + 125.94(0.83)(-0.18985) = -5.61$, which gives us a stock price at expiration of $125.94 - 5.61 = 120.33$, leading to an option price of $Max(0, 120.33 - 125) = 0$. We repeat this procedure several thousand times, after which we take an average of the simulated option prices and then discount that average to the present using the present value formula $e^{-0.0446(0.0959)}$.

Naturally every simulation is different because each set of random numbers is different. A Monte Carlo procedure written in Excel's Visual Basic produced the following values for this call, whose actual Black-Scholes-Merton price is 13.55:

Number of Random Drawings	Call Price
1,000	12.64
10,000	13.73
50,000	13.69
100,000	13.67

It would appear that at least 100,000 random drawings are required for the simplest case of a standard European option.

Applying the Monte Carlo technique to more complicated options such as path-dependent options requires a partitioning of the option's life into time periods, as in the binomial model. For example, suppose you wanted to price an Asian call option in which the average price would be computed by collecting the daily closing price over the life of the option. Ignoring holidays and weekends, let us say that the option has a 90-day life. Then a run would consist of 90 random drawings, each used to simulate the stock price at the end of each of the 90 days. The formula for each ΔS would be based on the previous day's closing price. The value of Δt would be 1/365. Then the average of the 90 stock prices would determine the call payoff at expiration as described in this chapter. You would then need to repeat the procedure a large number of times.

This may seem like a formidable task but that is what computers are for. More than likely computations such as these would be written in a fast and efficient language like C++. In addition, there is a considerable amount of research going on for ways to make Monte Carlo simulations run more efficiently.

For your purposes here, the important thing is to gain an understanding of the principles of option pricing with Monte Carlo simulation. For example, consider a Monte Carlo simulation of a European option. Each run generates a possible outcome. Provided enough runs are made, the outcomes will occur with the same relative frequency implied by the probabilities assumed by the Black-Scholes-Merton model. The option price will then become what we have so often described throughout this book—a probability-weighted average of the expiration values of the option, discounted at the risk-free rate. More complex options will naturally require modifications to the procedure.

15

FINANCIAL RISK MANAGEMENT TECHNIQUES AND APPLICATIONS

In the first 14 chapters we studied the use of many different types of derivatives in a variety of situations. In the financial markets in recent years, derivatives have been playing a major part in the decision-making process of corporations, financial institutions, and investment funds. Derivatives have been embraced not only as tools for hedging but as means of controlling risk; that is, reducing risk when one wants to reduce risk and increasing risk when one wants to increase it. The low transaction costs and the ease of using derivatives have given firms flexibility to make adjustments to the risk of a firm or portfolio. Corporations have been particularly avid users of derivatives for managing interest rate and foreign exchange risks.

As we have seen throughout this book, derivatives generally carry a high degree of leverage. When used improperly they can increase the risk dramatically, sometimes putting the survival of a firm in jeopardy. In fact, in Chapter 16, we shall take a look at a few stories of how derivatives were used improperly, and in some cases were fatal.

The critical importance of using derivatives properly has created a whole new activity called risk management. Risk management is the practice of defining the risk level a firm desires, identifying the risk level a firm currently has, and using derivatives or other financial instruments to adjust the actual level of risk to the desired level of risk. Risk management has also spawned an entirely new industry of financial institutions that offer to take positions in derivatives opposite the end users, which are corporations or investment funds. These financial institutions, which we have previously identified as dealers, profit off of the spread between their buying and selling prices and generally hedge the underlying risks of their portfolios of derivatives.

WHY PRACTICE RISK MANAGEMENT?

Impetus for Risk Management

Growth in the use of derivatives for managing risk did not occur simply because people became enamored with them. In fact, there has always been a great deal of suspicion, distrust, and outright fear of derivatives. Eventually, firms began to realize that derivatives were the best tool for coping with markets that had become increasingly volatile and over which most businesses felt were beyond their means to forecast and control.

The primary sources of these risks are interest rates, exchange rates, commodity prices, and stock prices. These are risks over which most businesses have little expertise. Obviously businesses must take some risks or there would be no reason to be in business. Acceptable risks are those related to the industries and products in which a business operates. These risks, which are often called strategic risks, are those in which a business should have some expertise. Risks driven by factors external to an industry or products, such as interest rates, exchanges rates, commodity prices, and stock prices are those over which a typical business has little strategic advantage. It makes sense, therefore, for a business to manage and largely eliminate these risks.

For example, consider an airline. Its strategic expertise is in transporting people safely from one destination to another. Yet, airlines assume a number of risks over which they have little control or expertise. The cost of fuel, borrowing costs, and exchange rates exert an enormous influence on airlines' performance. While some airlines merely accept all of these risks as a part of doing business, others choose to actively manage these risks. An increasing number of businesses have begun to recognize the benefits of this strategy of accepting risks over which they have some control and expertise while actively managing the other risks.

Another reason why many firms have begun to practice risk management is simply that they have learned a lesson by watching other firms. Seeing other companies fail to practice risk management, and watching them suffer painful and embarrassing lessons—or hearing of another company that has successfully developed a risk management system—can be a powerful motivator.

Naturally derivatives have emerged as a popular tool for managing risks, and companies have shown an increasing tendency to use derivatives. Financial institutions have made this growth possible by creating an environment that is conducive to the efficient use of derivatives. This environment depended heavily on the explosion in information technology witnessed in the 1980s and 1990s. Without enormous developments in computing power, it would not have been possible to do the numerous and complex calculations necessary for pricing derivatives quickly and efficiently and for keeping track of positions taken.[1]

Another factor that has fueled the growth of derivatives was the favorable regulatory environment. In the U.S., the CFTC adopted a pro-market position in the early 1980s, which paved the way for an increasing number of innovative futures contracts, such as Eurodollars and stock index futures. These contracts established a momentum that led to more innovation in the global exchange-listed and over-the-counter markets. Perhaps one of the most important steps taken by the CFTC was a non-step: its decision not to regulate over-the-counter derivatives transactions.

We should note that the derivatives industry has gone through a significant evolution in what it calls itself. In its early stages, it was known as *commodities.* As exchange-listed options and financial futures were created, it began to call itself *futures and options.* When over-the-counter products like swaps and forwards were added, it began to be known as *derivatives.* Now the focus has shifted away from the instruments and toward the process, leaving us with the term *risk management* and sometimes financial risk management.

Benefits of Risk Management

In Chapter 11, we identified several reasons why firms hedge. At this point, however, we are focusing not on the simple process of hedging, but on the more general process of managing risk. Let us restate our Chapter 11 reasons for hedging in the context of risk management.

In the Modigliani-Miller world in which there are no taxes or transaction costs and information is costless and available to everyone, financial decisions have no relevance for shareholders. Financial decisions, such as how much debt a firm issues, how large a dividend it pays, or how much risk it takes, merely determine how the pie is sliced. The size of the pie, as determined by the quality of a firm's investments in its assets, is what determines shareholder value. Modigliani and Miller argue that shareholders can do these financial transactions just as well by buying and selling stocks and bonds in their personal portfolios. Risk management is also a financial decision. Thus, risk management can, in theory, be practiced by shareholders by adjusting their personal portfolios; consequently, there is no need for firms to practice risk management.

---

[1]In fact, one could argue that the development of the personal computer was one of the single most important events for derivatives.

This argument ignores the fact that most firms can practice risk management more effectively and at lower cost than shareholders. Their size and investment in information systems give firms an advantage over their shareholders. Firms can also gain from managing risk if their income fluctuates across numerous tax brackets. With a progressive tax system, they will end up with lower taxes by stabilizing their income. Risk management can also reduce the probability of bankruptcy, a costly process in which the legal system becomes a partial claimant on the firm's value. Risk management can also be done because managers, whose wealth is heavily tied to the firm's performance, are simply managing their own risk.

Firms that are in a near-bankrupt state will find that they have little incentive to invest in seemingly attractive projects that will merely help their creditors by increasing the chance that the firm will be able to pay off its debts. This is called the underinvestment problem and is more thoroughly explored in corporate finance books. Managing risk helps avoid getting into situations like that and, as such, increases the chance that firms will always invest in attractive projects, which is good for society as a whole. Risk management also allows firms to generate the cash flow necessary to carry out their investment projects. If internal funds are insufficient, they may have to look toward external funds. Some firms would cut investment rather than raise new capital.

As described in an earlier section, when a firm goes into a particular line of business, it knowingly accepts risks. Airlines, for example, accept the risks of competition in the market for transporting people from one place to another. The risk associated with volatile oil prices is an entirely different type of risk, one that they often prefer to eliminate. Hence many airlines hedge oil prices, which allows them to concentrate on their main line of business. On occasion, however, they may feel that oil prices are heading downward, suggesting that they lift their hedges. Thus, they are not just hedging but rather practicing risk management by setting the current level of risk to the desired level of risk.

Some firms use risk management as an excuse to speculate in areas where they have less expertise than they think. As we shall see in Chapter 16, when a consumer products company speculates on foreign interest rates, it is no longer just a consumer products company. It becomes a financial trading company and it must be prepared to suffer the consequences if its forecasts are wrong. Some firms practice risk management because they truly believe that they can time movements in the underlying source of a risk. When that source of risk is unrelated to the firm's basic line of business, the consequences of bad forecasting combined with highly leveraged derivatives can be dire.

Other firms manage risk because they believe that arbitrage opportunities are possible. For example, suppose a firm could borrow at a floating rate of LIBOR plus 110 basis points or at a fixed rate of 10.5 percent. It can enter into a swap paying a fixed rate of 9.25 percent and receiving LIBOR. Simple arithmetic shows that it would get a lower rate by issuing floating-rate debt and entering into a pay fixed-receive floating swap. The net effect is to pay LIBOR plus 110 on the floating-rate debt, receive LIBOR, and pay 9.25 on the swap, which adds up to a fixed rate of $9.25 + 1.10 = 10.35$, or 15 basis points cheaper than straight fixed-rate debt. Yet if the firm had simply issued fixed-rate debt at 10.5 percent, it would have assumed no credit risk. Now it assumes the credit risk of the swap counterparty and is compensated to the tune of a 15 basis point reduction in the interest rate. Is this worth it? In the early days of the market, the savings were probably large enough to be worth it, but as the market has become more efficient, the savings have decreased and are likely just fair compensation for the assumption of credit risk. Nonetheless, the credit risk may be worth taking in order to lower borrowing costs.

It is important to emphasize that reducing risk is not in and of itself a sufficient reason to hedge or manage risk. Firms that accept lower risks will in the long run earn lower returns. Moreover, if their shareholders truly wanted lower risks, they could easily realign their portfolios, substituting lower-risk securities for higher-risk securities. Managing risk must create value for shareholders, giving them something they cannot get themselves. To the extent that risk management reduces the costly process of bankruptcy, saves taxes, and makes it easier for firms to take on profitable investment projects, value is clearly created.

In the next section we take a close look at how to manage the most important type of risk: market risk.

MANAGING MARKET RISK

Market risk is the uncertainty of a firm's value or cash flow that is associated with movements in an underlying source of risk. For example, a firm might be concerned about movements in interest rates, foreign exchange rates, stock prices, or commodity prices.

When considering interest rate risk, there is the risk of short-, intermediate-, and long-term interest rates. Within short-term interest rate risk, there is the risk of LIBOR changing, the risk of the Treasury bill rate changing, the risk of the commercial paper rate changing, and numerous other risks associated with specific interest rates. A risk manager responsible for positions in LIBOR-based instruments and instruments tied to the commercial paper rate would have to take into account the extent to which those rates are correlated. A long position in LIBOR and a short position in commercial paper would be a partial hedge since LIBOR is correlated with the commercial paper rate. Thus, the combined effects of all sources of risk must be considered.

The effects of changes in the underlying source of risk will show up in movements in the values of spot and derivative positions. You should recall that in Chapter 5 we introduced the concept of an option's delta, which was the change in the option's price divided by the change in the underlying stock's price. We noted that a delta-hedged option would move perfectly with and be offset by an appropriately weighted position in the stock. A delta-hedged portfolio would be neutral with respect to stock price movements, but only for very small stock price changes. For large stock price changes, the delta may move too quickly. We noted that the risk of the delta changing too quickly is captured by the option's gamma. We also saw that if the volatility of the underlying stock changes, the option price can change quite significantly, even without a movement in the stock price. This risk is captured by the option's vega. These delta, gamma, and vega risk measures are equally applicable to many instruments other than options and stocks. They are some of the tools used by risk managers to control market risk.

Let us consider a situation in which, to accommodate a customer, a derivatives dealer has taken a position in a $10 million notional principal four-year interest rate swap that pays a fixed rate and receives a floating rate. In addition, the dealer has sold a three-year $8 million notional principal interest rate call with an exercise rate of 12 percent. We assume that for both instruments the underlying is LIBOR. To keep things as simple as possible, we shall let the payments on the swap occur once a year. The current term structure of LIBOR and the implied forward rates are shown at the top of Table 15.1.

Let us first price the interest rate swap. Using the procedure we learned in Chapter 12, we see in Table 15.1 that the rate is 11.85 percent. Now let us price the three-year interest rate call. Table 15.1 uses the Black model we learned about in Chapter 13 and shows that the premium on this option would be $73,745. Now let us look at how to delta hedge this combination of a swap and an option.

Delta Hedging

In order to delta hedge, we must make the portfolio be unaffected by small movements in interest rates. To do this we shall need the delta of the swap and the option. Either can be obtained by taking the mathematical first derivative of the swap or option value with respect to interest rates. In this example, however, we shall estimate the delta by repricing both instruments when the one-period spot rate and the remaining forward rates move up and down one basis point. Then we shall average the movement in the derivative's price, which will be a good approximation of the delta.[2]

If all forward rates move up one basis point, the new one-period spot rate will be 10.01 percent and the new forward rates will be 12.02 percent, 12.82 percent, and 13.22 percent. This necessitates recalculating the spot rates, which will be 11.01 percent, 11.61 percent, and 12.01 percent.[3] Using the procedure we learned in

[2]The price change for a one basis point move up is slightly different from that for a one basis point move down. This is why we take the average price change. This effect results from the convexity of the price curve and plays a role in gamma hedging.

[3]It appears as if the new spot rates are just one basis point above the old spot rates. This is not precisely the case, as would be indicated if we let the forward rates shift by a much larger amount or if we carried our results out to more significant digits.

Recall that we are short the option so gains occur on rate decreases and losses occur on rate increases.

As we learned in Chapter 5, the risk associated with larger price moves in which the delta does not fully capture the risk is called gamma risk. It is the risk of the delta changing. To be fully hedged a dealer would have to be delta hedged at all times. If rates move sharply, the effective delta would not equal the actual delta until the dealer could put on another transaction that would reset its delta to the appropriate value. This could be too late. This risk can be hedged, however, by combining transactions so that the delta and gamma are both zero. First, however, we must estimate the gamma. Table 15.3 illustrates the calculation of the gammas of the swap and the option. The gammas are −$12,500 for the swap and $5,000 for the option. This means that as LIBOR increases, the swap delta decreases in value by $12,500(0.0001) = $1.25 and the option delta increases in value by $5,000(0.0001) = $0.50. Because we are short the option, its gamma is actually −$5,000. Thus, our overall gamma is −$17,500.

Assuming that we have delta hedged the swap and option with the Eurodollar futures, our gamma will still be −$17,500 because the gamma of the futures is zero. We are delta hedged but not gamma hedged. To

Table 15.3 Estimation of Swap and Option Gammas

Basis Point Change	Swap Value	Average Change in Swap Value[1]	Swap Gamma[2]	Option Value	Average Change in Option Value[3]	Option Gamma[4]
−0.0002	−$4,263			$73,258		
−0.0001	−2,131	$2,131.50		73,501	$243.50	
0.0000	0	2,130.50	−$12,500	73,745	244.00	$5,000
+0.0001	2,130	2,129.00		73,989	244.50	
+0.0002	4,258			74,234		

[1]The average change in the swap value is estimated as follows:

From a basis point chance of 0.0000: $\frac{(0-(-2{,}131))+2{,}130-0}{2}=22{,}130.50.$

From a basis point chance of 0.0001: $\frac{(2{,}130-0)+(4{,}258-2{,}130)}{2}=2{,}129.00.$

From a basis point chance of −0.0001: $\frac{(-2{,}131-(-4{,}263)+(0-(-2{,}131)}{2}=2{,}131.50.$

These calculations are the deltas at these points.

[2]The swap gamma is estimated as follows: $\frac{(-2{,}130.50-2{,}130.50)+(2{,}129.00-2{,}130.50)}{2}=-1.25.$

A change in the delta of −1.25 for a one basis point move implies a gamma of −1.25/0.0001 = −$12,500.

[3]The average change in the option value is estimated as follows:

From a basis point chance of 0.0000: $\frac{(73{,}745-73{,}501)+(73{,}989-73{,}745)}{2}=244.00.$

From a basis point chance of 0.0001: $\frac{(73{,}989-73{,}745)+(74{,}234-73{,}989)}{2}=244.50.$

From a basis point chance of −0.0001: $\frac{(73{,}501-73{,}258)+(73{,}745-73{,}501)}{2}=243.50.$

[4]The option gamma is estimated as follows: $\frac{(244.00-243.50)+(244.50-244{,}00)}{2}=0.50.$

A change in the delta of 0.50 for a one basis point move is a gamma of 0.50/0.0001 = $5,000.

become gamma hedged we will need another instrument. Let us assume that the instrument chosen is a one-year call option with an exercise rate of 11 percent whose delta is $43 and whose gamma is $2,500, both figures under the assumption of a $1 million notional principal. The problem is to determine the appropriate notional principal of this new option so that we will be delta hedged and gamma hedged. This is a simple problem answered by solving simultaneous equations.

Let us assume that we take x_1 Eurodollar futures, which have a delta of −$25 and a gamma of zero; and x_2 of the one-year calls, which have a delta of $43 and a gamma of $2,500 per $1,000,000 notional principal. Our swap and other option have a delta of $1,887 and a gamma of −$17,500. We eliminate the delta and gamma risk of the portfolio by setting the delta to zero by the equation

$$\$1{,}887 + x_1(-\$25) + x_2(\$43) = \$0 \quad \text{(Delta)},$$

and the gamma to zero by the equation

$$-\$17{,}500 + x_1(\$0) + x_2(\$2{,}500) = \$0 \quad \text{(Gamma)}.$$

These are simply two equations with two unknowns. The solution is an exercise in basic algebra, but just to make sure you understand, we shall work through it. Rewrite the equations as

$$x_1(-\$25) + x_2(\$43) = -\$1{,}887$$
$$x_2(\$2{,}500) = \$17{,}500.$$

Solve the second equation for $x_2 = 7.00$. Then insert 7.00 for x_2 in the first equation and solve for x_1 to get $x_1 = 87.52$. This means we need to go long 87.52 Eurodollar futures. Round to 88. The solution $x_2 = 7.00$ means that we need to go long 7.00 times the notional principal of $1,000,000 on which the new option's delta and gamma were calculated. In other words, we need $7,000,000 notional principal of the one-year option. These transactions combine to set the delta and gamma of the overall position to approximately zero.[5]

Unfortunately the use of options introduces a risk associated with possible changes in volatility. Let us take a look at how that risk arises and how we can hedge it.

Vega Hedging

In Chapter 5 we learned that the change in the option price over the change in its volatility is called its vega. A portfolio of derivatives that is both delta and gamma hedged can incur a gain or loss even when there is no change in the underlying as a result of a change in the volatility. Most options are highly sensitive to the volatility, which changes often. Consequently it is important to try to hedge vega risk.

Swaps, futures, and FRAs do not have vegas because volatility is not a determinant of their prices. In our example, we need consider only the vega of the three-year call option. Although option pricing formulas can often give the vega in its exact mathematical form, we shall estimate the vega by changing LIBOR by one basis point in each direction. Recall that under the initial term structure, the option value is $73,745. If the volatility increases from 0.147 to 0.1471, the new option value will be $73,787, a change of $42. If volatility decreases from 0.147 to 0.1469, the new option value will be $73,703, a change of −$42. This is an average change of $42. Since we are short this option, the vega of our portfolio of the four-year swap and this three-year option is −$42.

Now consider the one-year option that we introduced in the last section to use for gamma hedging. It will also have a vega that must be taken into account. Its vega is estimated to be $3.50 for every $1,000,000 notional principal. Our delta- and gamma-hedged portfolio would have $7 million face value of this option, making the vega $24.50. That would make our overall portfolio have a vega of $24.50 − $42, or −$17.50. Thus, we still have a significant risk that volatility will increase, and each 0.0001 increase in volatility will cost us $17.50.

---

[5]As a check, we see that the delta is 1,887 + 88(−25) + 7(43) = −12.00 and the gamma is −17,500 +7(2,500) = 0.

In order to hedge delta, gamma, and vega we will need three hedging instruments. Because of the vega risk, at least one of the instruments has to be an option. Let us use an option on a Eurodollar futures that trades at the Chicago Mercantile Exchange alongside the Eurodollar futures. The option has a delta of −\$12.75, a gamma of −\$500, and a vega of \$2.50 per \$1,000,000 notional principal. This leads to the following set of simultaneous equations:

$$\$1{,}887 + x_1(-\$25) + x_2(\$43) + x_3(-\$12.75) = 0 \text{ (Delta)}$$
$$-\$17{,}500 + x_1(\$0) + x_2(\$2{,}500) + x_3(-\$500) = 0 \text{ (Gamma)}$$
$$-\$42 + x_1(\$0) + x_2(\$3.50) + x_3(\$2.50) = 0 \text{ (Vega).}$$

The first equation sets the portfolio delta to zero, the second sets the gamma to zero, and the third sets the vega to zero. The coefficients x_1, x_2, and x_3 represent quantities of \$1,000,000 notional principal that should be established with Eurodollar futures, the one-year option, and the option on the Eurodollar futures. To solve these equations, we first note that the second and third equations can be written as

$$x_2(\$2{,}400) + x_3(-\$500) = \$17{,}500$$
$$x_2(\$3.50) + x_3(\$2.50) = \$42,$$

which is simply two equations with two unknowns. Multiplying the second equation by 200 gives us $\$700x_2 + \$500x_3 = \$8{,}400$. Then adding the two equations gives $\$3{,}200x_2 = \$25{,}900$, which gives $x_2 =$ 8.09375. Inserting 8.09375 for x_2 in either of these equations gives $x_3 = 5.46875$. Thus, we need 8.09375(\$1,000,000) = \$8,093,750 notional principal of the four-year option and 5.46875(\$1,000,000) = \$5,468,750 notional principal of the Eurodollar futures option. We then insert 8.09375 and 5.46875 into the first equation for x_2 and x_3, giving us $x_1(-\$25)$ 18.09375(\$43) + 5.46875(−\$12.75) = −\$1,887. Solving for x_1 gives a value of 86.61. This means that we would buy 87 Eurodollar futures.[6]

It should be apparent by now that the dealer should not hedge by setting the delta to zero and then attempting to hedge the gamma and vega risk with other instruments. As these instruments are added to eliminate gamma and vega risk, the delta hedge is destroyed. There are two possible approaches to solving the problem, one being the simultaneous equation approach that we followed here. It is guaranteed to provide the correct solutions. Another approach would be to solve the gamma and vega hedge simultaneously, which will set the gamma and vega to zero, but leave the overall delta nonzero. Then the delta hedge can be set with Eurodollar futures, which have a delta but no gamma or vega. Consequently, adding them to the position at the very end will not change the gamma or vega neutrality.

In spite of a dealer's efforts at achieving a delta-gamma-vega neutral position, it is really impossible to achieve an absolute perfect hedge. The vega hedge is accurate only for extremely small changes in volatility. Large changes would require yet another adjustment. In addition, all deltas, gammas, and vegas are valid only over the next instant in time. Even if there were no changes in LIBOR or the volatility, the position would become unhedged over time if no further adjustments were made. Eventually the portfolio would become significantly unhedged, so some adjustments might be made to realign the portfolio to a delta-gamma-vega neutral position, possibly as often as once a day.

It is apt to remember a famous expression: *The only perfect hedge is in a Japanese garden.* Any dealer accepts the fact that a small amount of risk will be assumed. To date, however, no major derivatives dealer who has made the effort to be hedged has suffered a significant loss and most have found market making in derivatives to be a moderately profitable activity with very low risk. This is a testament to the excellent risk management practiced by the major derivatives dealers.

---

[6]Because of rounding off to whole numbers of contracts, the overall position is not quite perfectly hedged. The delta is \$1,887 + 87(−\$25) + 8.09375(\$43) + 5(−\$12.75) = \$3.72. The gamma is −\$17,500 + 87(\$0) + 8.09375(\$2,500) + 5(−\$500) = \$234,38. The vega is −\$42 + 8.09375(\$3.50) + 5(\$2.50) = −\$1.17.

On the other side of the transaction is the end user, the party who approaches the dealer about entering into a derivatives transaction. Most end users are corporations attempting to hedge their interest rate, currency, equity, or commodity price risk. Some will speculate from time to time. Most, however, already have a transaction in place that has a certain amount of risk. They contract with the dealer to lay off that risk. Rarely will the end user engage in the type of dynamic hedging illustrated above. That is because the end user is not typically a financial institution like the dealer. Financial institutions can nearly always execute transactions at lower cost and can generally afford the investment in expensive personnel, equipment, and software necessary to do dynamic hedging. Most end users enter into derivatives transactions that require little or no adjustments. You have, of course, seen many such examples throughout this book. Some end users have, however, suffered losses from being unhedged at the wrong time or from outright speculating. Yet, most end users could have obtained a better understanding about the magnitude of their risk and the potential for large losses had they applied the technique called Value at Risk or VAR.

Value at Risk (VAR)

Value at Risk, or VAR, is a dollar measure of the minimum loss that would be expected over a period of time with a given probability. For example, a VAR of $1 million for one day at a 5 percent probability means that the firm would expect to lose at least $1 million in one day 5 percent of the time. Some prefer to express such a VAR as a 95 percent probability that a loss will not exceed $1 million. In this manner, the VAR becomes a maximum loss with a given confidence level. The significance of a million dollar loss depends on the size of the firm and its aversion to risk. But one thing is clear from this probability statement: a loss of at least $1 million would be expected to occur once every 20 trading days, which is about once per month.

VAR is widely used by dealers, even though their hedging programs nearly always leave them with little exposure to the market. If dealers feel that it is important to use VAR, that should be a good enough reason for end users to employ it, and surveys show that an increasing number of end users are doing so.

The basic idea behind VAR is to determine the probability distribution of the underlying source of risk and to isolate the worst given percentage of outcomes. Using 5 percent as the critical percentage, VAR will determine the 5 percent of outcomes that are the worst. The performance at the 5 percent mark is the VAR.

Table 15.4 provides a simple illustration with a discrete classification of the change in the value of a hypothetical portfolio. Note that each range has a probability and a cumulative probability associated with it. Starting with the class with the worst outcome, VAR is found by examining the cumulative probability until the specified percentage is reached. In this case, VAR for 5 percent is $3,000,000. This would be interpreted as follows: There is a 5 percent probability that over the given time period, the portfolio will lose at least $3 million. Of course VAR can be expressed with respect to any chosen probability. Such statements as "There is a 15 percent probability that over the given time period, the portfolio will lose at least $2 million" and "There is a 50 percent probability that over the given time period, the portfolio will incur a loss" are both legitimate statements of the portfolio's VAR.

Table 15.4 Probability Distribution of Change in Portfolio Value

Change in Portfolio Value	Probability	Cumulative Probability
−$3,000,000 and lower	0.05	0.05
−$2,000,000 to −$2,999,999	0.10	0.15
−$1,000,000 to −$1,999,999	0.15	0.30
$0 to −$999,999	0.20	0.50
$0 to $999,999	0.20	0.70
$1,000,000 to $1,999,999	0.15	0.85
$2,000,000 to $2,999,999	0.10	0.95
$3,000,000 and higher	0.05	1.00

Figure 15.1 illustrates the principle behind VAR when the distribution of the portfolio change in value is continuous. The familiar normal or bell-shaped curve is widely used though not necessarily appropriate in many cases. Accepting it as legitimate for our purposes, we see where the 5 percent VAR is noted, which is 1.65 standard deviations from the expected change in portfolio value, which in this example the expected change is zero. Of course, not all portfolios have an expected change of zero. In any case, the rule for determining VAR when applying normal probability theory is to move 1.65 standard deviations below the expected value. Beyond that point, 5 percent of the population of possible outcomes is found. For a VAR of 1 percent, you would move 2.33 standard deviations below the expected value.

FIGURE 15.1 Value at Risk for Normally Distributed Change in Portfolio Value with Zero Expected Change

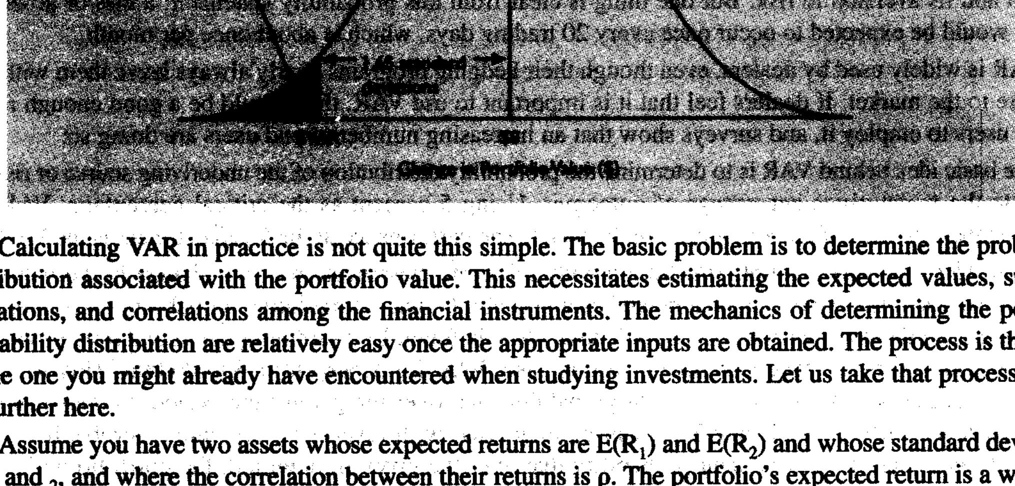

Calculating VAR in practice is not quite this simple. The basic problem is to determine the probability distribution associated with the portfolio value. This necessitates estimating the expected values, standard deviations, and correlations among the financial instruments. The mechanics of determining the portfolio probability distribution are relatively easy once the appropriate inputs are obtained. The process is the same as the one you might already have encountered when studying investments. Let us take that process a little bit further here.

Assume you have two assets whose expected returns are $E(R_1)$ and $E(R_2)$ and whose standard deviations are $_1$ and $_2$, and where the correlation between their returns is ρ. The portfolio's expected return is a weighted average of the expected returns of assets 1 and 2,

$$E(R_p) = w_1E(R_1) + w_2E(R_2),$$

where w_1 and w_2 are the percentages of the investor's wealth that are allocated to assets 1 and 2, respectively. The portfolio standard deviation is a more complicated weighted average of the variances of assets 1 and 2 and their covariance,

$$\sigma_p = \sqrt{w_1^2\sigma_1^2 + w_1^2\sigma_2^2 + 2w_1w_2\sigma_1\sigma_{2\rho}},$$

where the expression $_{12}$ is recognized as the covariance between assets 1 and 2.

There are three methods of estimating VAR.

Analytical Method The analytical method, also called the variance-covariance method, makes use of knowledge of the input values and any necessary pricing models along with an assumption of a normal distribution.

We illustrate the analytical method with two examples. Suppose a portfolio manager holds two distinct classes of stocks. The first class, worth \$20 million, is identical to the S&P 500. It has an expected return of 12 percent and a standard deviation of 15 percent. The second class is identical to the Nikkei 300, an index of Japanese stocks, and is valued at \$12 million. We shall assume the currency risk is hedged. The expected return is 10.5 percent and the standard deviation is 18 percent. The correlation between the Nikkei 300 and the S&P 500 is 0.55. All figures are annualized. With this information we can calculate the portfolio expected return and standard deviation as

$$E(R_p) = (20/32)0.12 + (12/32)0.105 = 0.1144$$

$$\sigma_p = \sqrt{(20/32)^2(0.15)^2 + (12/32)^2(0.18)^2 + 2(20/32)(12/32)(0.15)(0.18)(0.55)}$$

$$= 0.1425.$$

Now let us calculate this portfolio's VAR at a 5 percent level for one week. First we must convert the annualized expected return and standard deviation to weekly equivalents. This is done by dividing the expected return by 52 (for the number of weeks in a year) and dividing the standard deviation by the square root of 52, which is 7.21. This gives us 0.1144/52 = 0.0022 and 0.1425/7.21 = 0.0198. Under the assumption of a normal distribution, the return that is 1.65 standard deviations below the expected return is

$$0.0022 - 1.65(0.0198) = -0.0305.$$

The portfolio would be expected to lose at least 3.05 percent 5 percent of the time. VAR is always expressed in dollars, so the VAR is \$32,000,000(0.0305) = \$976,000. In other words, the portfolio would be expected to lose at least \$976,000 in one week 5 percent of the time or one out of twenty weeks.

Let us now calculate VAR for a portfolio containing options. In fact, let us make it an extremely risky portfolio, one consisting of a short call on a stock index. We assume that the call has one month to go before expiring, the index is at 720, the exercise price is 720, the risk-free rate is 5.8 percent, and the volatility of the index is 15 percent. We shall ignore dividends. Inserting these figures into the Black-Scholes-Merton model tells us that the call should be priced at \$14.21. Assume that the index option contract has a multiplier of 500, so the total cost is 500(\$14.21) = \$7,105. We shall assume that the investor sells 200 contracts, resulting in the receipt up front of 200(\$7,105) = \$1,421,000. The worst outcome for an uncovered call is for the stock to increase. Let us look at the 5 percent worst outcomes, which occur on the upside. Using the same information on the index from the previous example, we must first convert to monthly data. The expected return on the index is 0.1144/12 = 0.0095 and 0.1425/3.46 (the square root of 12) = 0.0412. On the upside the 5 percent tail of the distribution is 0.0095 + 1.65(0.0412) = 0.0775. That would leave the index at 720(1.0775) = 775.80 or higher. If the option expires with the index at 775.80, it will have a value of 775.80 − 720 = 55.80. Thus, our net loss will be 55.80 − 14.21 = 41.59 per option. The total loss will be 200(500)(41.59) = \$4,159,000. Thus, the VAR for this short call is \$4,159,000 and we can, therefore, say that the portfolio will lose at least \$4,159,000 in one month 5 percent of the time. This would be once every 20 months.

Although we calculated an expected value in these examples, it is fairly common to assume a zero expected value. This is because one day is a common period over which to calculate a VAR and the expected daily return is very small. A typical VAR calculation is much more highly influenced by the volatility than by the expected return.

The analytical method uses knowledge of the parameters of the probability distribution of the underlying sources of risk at the portfolio level. Since the expected value and variance are the only two parameters used, the method implicitly is based on the assumption of a normal distribution. If the portfolio contains options, the assumption of a normal distribution is no longer valid. Option returns are highly skewed and the expected return and variance of an option position will not accurately produce the desired result, the return that is exceeded, say, 5 percent of the time. One approach is the one used here: We identified the critical outcome of the underlying and then determined the option outcome that corresponds to it.

Another commonly used alternative employs the delta, rather than the precise option pricing model, to determine the option outcome. In fact, the analytical or variance-covariance method is also sometimes called the delta normal method. Although this method is only approximate, it has some advantages. The delta is a linear adjustment of the underlying price change to the option price change and linearity is a desirable and simplifying property. When the outcome of a normal distribution is adjusted in a linear manner, the result remains normally distributed. Thus, the delta normal approach linearizes the option distribution; in other words, it converts the option's distribution to a normal distribution. This can be useful, particularly when a large portfolio is concerned. For longer periods, such as the one-month period used here, the delta adjustment is sometimes supplemented with a gamma adjustment.

Another important concern in using the analytical method is that large portfolios can be very complicated to work with. In these examples, we identified only a single source of risk. For large institutions, there are literally thousands of sources of risk. The volatilities and correlations of these diverse sources of risk must be captured and consolidated into a single volatility for the portfolio. This requires massive amounts of information. Fortunately, this information is readily available. The RiskMetrics Corporation, http://www.riskmetrics.com, a spin-off of the noted Wall Street firm J. P. Morgan, has provided downloadable data sets, which are updated daily, on the Internet. This information is based on recent historical price behavior and is smoothed via a weighting of current and past volatility.

The primary advantage and disadvantage of the analytical method is its reliance on the assumption of a normal distribution. The following method gets around that assumption.

Historical Method The historical method estimates the distribution of the portfolio's performance by collecting data on the past performance of the portfolio and using it to estimate the future probability distribution. Obviously it assumes that the past distribution is a good estimate of the future distribution.

Figure 15.2 presents an example of the information obtained from a historical sample of daily portfolio returns over approximately one year, 254 trading days in this case. The information is presented in the form of a histogram. The person doing the analysis would choose the most appropriate intervals. In this case, the 13 worst outcomes, which is about 5.1 percent of the total possible, are shaded. This means that the critical portfolio VAR is a loss of 5.1 percent. If the portfolio size is \$15 million, then the VAR is \$15,000,000(0.10) = \$1,500,000. Thus, we can say that the portfolio would be expected to lose at least \$1.5 million in one day about 5 percent of the time, which is about once per month.

The historical method obviously produces a VAR that is consistent with the VAR of the chosen historical period. Whether this is an accurate method depends on several factors. Obviously it matters greatly whether the probability distribution of the past is repeated in the future. Also, the portfolio held in the future might differ in some way from the portfolio held in the past. For example, the S&P/Nikkei portfolio, which is currently \$20 million in the S&P and \$12 million in the Nikkei, could be reallocated. In fact, unless each asset performs identically, the portfolio is automatically reallocated. One asset grows in value at a greater rate than the other and this changes the portfolio's expected return and standard deviation. This problem, however, can be accommodated by using the historical returns but applying new weights to each asset in accordance with the current weights rather than the historical weights.

Another limitation of the historical method is that it requires the choice of a sample period. The outcome can be greatly affected by how large a sample one selects. Normally, the larger a sample, the more reliable the estimates obtained from it, but the larger the sample, the older are some of the data and the less reliable they become.

Another problem with the historical method is that the historical period may be badly representative of the future. For example, suppose the historical period included the stock market crash of October 1987, a day

FIGURE 15.2 Histogram of Portfolio Returns for 254 Days

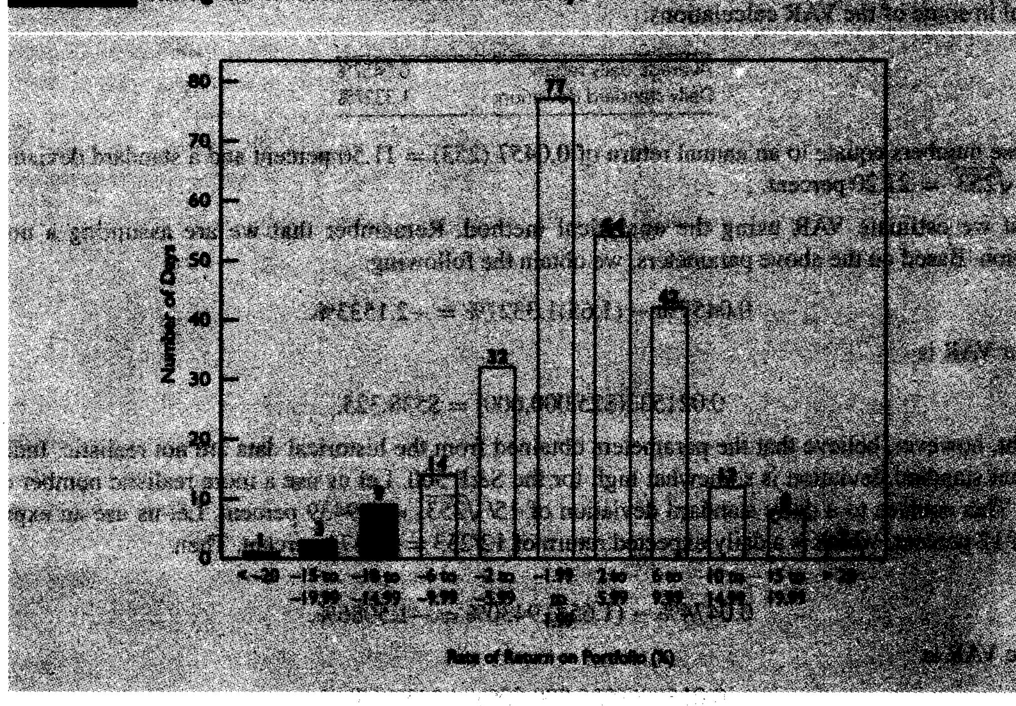

in which the market lost over 20 percent of its value in one day. Is this extreme outcome an accurate reflection of the VAR? Such an outcome would tremendously bias the volatility. Yet days like that are the very outcomes that risk managers should be worrying about.

The final method combines many of the best properties of the previous two methods.

Monte Carlo Simulation Method The Monte Carlo simulation method is based on the idea that portfolio returns can be fairly easily simulated. Simulation requires inputs on the expected returns, standard deviations, and correlations for each financial instrument. In Appendix 14 we examined one procedure for simulating stock prices called the Monte Carlo method. Essentially the same procedure is used when calculating VAR except that it is necessary to ensure that the portfolio's returns properly account for the correlations among the financial instruments. This means that one cannot simply independently generate returns for, say, the S&P 500 and Treasury bonds. One set of returns can be generated but the other set of returns must reflect any correlation between the two sets of returns. This is done by making an assumption about the process in which asset returns are generated.

Monte Carlo simulation is probably the most widely used method by sophisticated firms. It is the most flexible method, because it permits the user to assume any known probability distribution and can handle relatively complex portfolios; however, the more complex the portfolio, the more computational time required. Indeed Monte Carlo simulation is the most demanding method in terms of computer requirements. Nonetheless, the vast improvements in computing power in recent years have brought Monte Carlo to the forefront in risk management techniques.

A Comprehensive Calculation of VAR

Let us illustrate the three methods of calculating VAR by using a comprehensive example. Suppose we hold a very simple portfolio consisting of $25 million invested in the S&P 500. We would like to estimate VAR at 5 percent for one day using each of the three methods. First, we collect a sample of the daily returns on the S&P 500 for the past year. This gives us 253 returns. We obtain the following information, which will be useful in some of the VAR calculations:

Average daily return:	0.0457%
Daily standard deviation:	1.3327%

These numbers equate to an annual return of 0.0457 (253) = 11.56 percent and a standard deviation of $1.3327\sqrt{253} = 21.20$ percent.

First we estimate VAR using the analytical method. Remember that we are assuming a normal distribution. Based on the above parameters, we obtain the following:

$$0.0457\% - (1.65)1.3327\% = -2.1533\%.$$

Thus, our VAR is

$$0.021533(\$25{,}000{,}000) = \$538{,}325.$$

We might, however, believe that the parameters obtained from the historical data are not realistic. Indeed a 21 percent standard deviation is somewhat high for the S&P 500. Let us use a more realistic number of 15 percent. This equates to a daily standard deviation of $15/\sqrt{253} = 0.9439$ percent. Let us use an expected return of 12 percent, which is a daily expected return of 12/253 = 0.0474 percent. Then

$$0.0474\% - (1.65)0.9430\% = -1.5086\%.$$

Thus, the VAR is

$$0.015086(\$25{,}000{,}000) = \$377{,}150.$$

Which is more correct? That depends on which input seems more reasonable.

Remember that the analytical method is based on the assumption of a normal distribution. Is this a reasonable assumption? Figure 15.3 shows the distribution of returns for the S&P 500 that was used in this

FIGURE 15.3 [illegible]

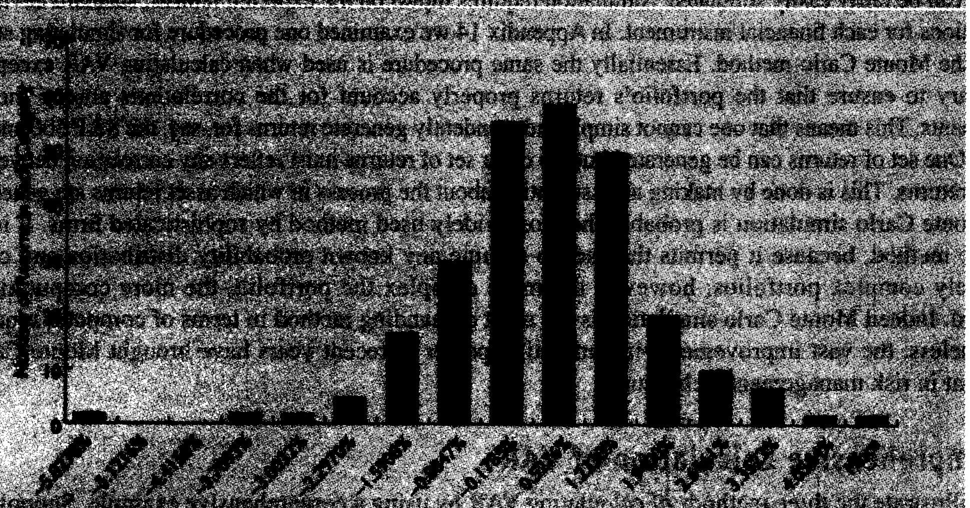

Note: [illegible]

example. We see that the distribution does indeed bear some resemblance to the normal, but it probably does not completely adhere to the normal distribution.

Now let us calculate VAR based on the historical method. This method simply requires that we rank the returns from worst to best. For 253 returns, we have 0.05(253) = 12.65, telling us that the VAR is the twelfth or thirteenth worst outcome. Let us use the thirteenth worst outcome. If we lined up the returns from worst to best, we would find that the thirteenth worst return would be −2.0969 percent. Thus, the VAR would be

$$0.020969(\$25{,}000{,}000) = \$524{,}225.$$

Now let us try a Monte Carlo simulation, where we have considerable flexibility. We must assume a probability distribution to generate the returns and we must choose the input values. Although one advantage of the Monte Carlo simulation is that we do not need to adhere to the normal distribution, we shall do so here for convenience and because it is the probability distribution that you are most likely to be very familiar with. In using the normal distribution we must input an expected value and volatility. Let us use the same 12 percent for the former and 15 percent for the latter that we used in the analytical method. We also have to decide on a time period, which we choose to be one day. In addition we must decide on the number of outcomes. Normally many thousands are used, but here we shall use 253, which is about one year and is the number of days over which we collected data. There is no particular reason why we should generate one year's worth of outcomes and, in practice, we would definitely want many more outcomes.

As noted earlier, the exact procedure for a Monte Carlo simulation with a normal distribution was described in the appendix to Chapter 14. Basically we generate a daily return according to the formula

[redacted]

where ε is a random number drawn from a standard normal distribution with a mean of zero and a standard deviation of 1.0. Microsoft Excel® will generate such numbers with its Analysis ToolPack. Alternatively, a good approximation can be obtained as the sum of twelve random numbers (rand() + rand() + · · · −6.0).

We do this 253 times, thereby obtaining 253 random returns. Once again, we simply sort them from worst to best and obtain the thirteenth worst return. Here, this is −1.3942 percent. Thus, the VAR is

$$0.013942(\$25{,}000{,}000) = \$348{,}550.$$

Of course, different values will be obtained from one simulation to another.

We now have a number of candidates for our VAR:

$538,325 (Analytical Method), or
$377,150 (Adjusted Analytical Method), or
$524,225 (Historical Method), or
$348,550 (Monte Carlo Simulation Method).

Which is correct? We simply do not know. The methods are so different and the inputs are so critical that this is a common result. In fact, even more variation could be obtained if we collected our data over different time periods or, for example, if we used weekly or monthly data instead of daily data. Or we could run more simulations or perhaps use a nonnormal distribution. The possibilities are almost endless. Moreover, the range of numbers shown here is actually quite small compared to what is frequently found in practice. Most real-world portfolios are far more complicated than this. They usually contain options, which tend to cause larger discrepancies. In short, VAR is a number that is quite sensitive to how it is calculated. That does not mean it is useless or unreliable. Knowing the potential range of VAR is itself useful. Moreover, one should always follow up the calculation with an ex post evaluation. If we settle on a VAR of, say, $400,000 in this example, we should expect that this value will be exceeded no more than 5 percent of the time. If it is exceeded far more or far less than 5 percent of the time, we know that $400,000 was not a good estimate of the VAR.

Assessing the accuracy of VAR systems is an important component of maintaining a quality enterprise risk system. Presently there are many different procedures for making this assessment, although they are quite technical.[7]

Benefits and Criticisms of VAR

Although widely criticized, VAR has been embraced by the risk management industry. VAR, or some variation thereof, is used by nearly every major derivatives dealer and an increasing number of end users. VAR is perhaps most beneficial in communicating information to nontechnical personnel. To tell a CEO that a firm is expected to lose at least $400,000 in a day 5 percent of the time, meaning about once every month, conveys a lot of useful information that the CEO can easily grasp. The trade-off, however, is that if the number is inaccurate, the CEO will have less confidence in the number and in the person giving him or her the number. Accurate estimation with continuing ex post follow-up is critical in using VAR.

VAR is also widely used in banking regulation. The objective of banking regulators is to ensure that the banking system does not fail and that consumers and savers are protected. Most banking regulators use VAR as a measure of the risk of a bank. A common specification for VAR in this context is 10 days with a 1 percent probability.

Likewise, banks and corporations who are engaged in significant trading activities commonly use VAR as a measure to allocate capital. In other words, they set aside a certain amount of capital to protect against losses. The amount of capital set aside is often the VAR.

VAR is also used in the evaluation of the performance of investment managers and traders. The modern approach to performance evaluation is to adjust the return performance for a measure of the risk taken in achieving that performance. VAR is often used in this context as a measure of risk.

Thus, we see that VAR has a variety of practical applications. Nonetheless, VAR must be used carefully. We have barely scratched the surface of VAR, but we have encountered enough to get you started in understanding it. There is much written on VAR; in fact, it is probably one of the most written-about topics of recent years. You will not be able to operate in the risk management world without encountering something written about it.

Extensions of VAR

In addition to estimating VAR, a risk manager will often subject the portfolio to a stress test, which determines how badly the portfolio will perform under some of the worst and most unusual circumstances. Consider the portfolio we discussed earlier that consisted of $20 million invested in the S&P 500 and $12 million invested in the Nikkei 300. Let us presume that in a given week both markets perform terribly, with the Nikkei losing 30 percent of its value and the S&P 500 losing 25 percent of its value. Then the Nikkei would lose $3.6 million and the S&P 500 would lose $5 million for a total loss of $8.6 million, which is 26.9 percent of the total value of the portfolio. Although such an outcome is extremely unlikely, a risk manager might test the portfolio's tolerance for such a remote possibility. If the performance is tolerable, then the portfolio risk is assumed to be acceptable. Stress testing can be quite valuable as a supplement to VAR and other techniques of risk management. Yet, stress testing has its own criticisms, including the fact that it places a tremendous emphasis on highly unlikely events.

The most important point to remember about VAR and stress testing, as well as any financial statistic, is that these are just estimates and cannot be expected to provide the full picture. Hence, some variations of VAR are widely used as a supplemental source of risk management information. Some risk managers focus on the expected loss beyond the VAR. For example, consider the VAR calculation we did for the S&P 500. Using the historical method, we ranked the outcomes and found that the VAR was 2.0969 percent, or

[7]See, for example, Joe H. Sullivan, Zachary G. Stoumbos, and Robert Brooks, "Real-time assessment of value-at-risk and volatility accuracy." *Nonlinear Analysis: Real World Applications*, 2007.

\$524,225. If we averaged all of the outcomes worse than −2.0969 percent, we would obtain an average loss of −2.8778 percent or \$719,450. This figure can be interpreted as the expected loss, given that the loss incurred exceeds the VAR. This figure is also sometimes known as the Conditional Value at Risk, because it reflects the expected loss conditional on the loss exceeding the VAR.

Value at Risk is a reasonable technique for use by firms that have assets whose values can be measured fairly easily. A financial institution would be one of those types of firms. Many corporations, however, have assets that generate cash flows but whose values cannot be easily determined. Consider for example, a company whose primary business is copper mining. Expenses are incurred in the mining process, and cash is generated to cover those expenses when the copper is sold. Of course in theory a copper mine has a market value, which would presumably come from discounting the stream of cash flows expected over the life of the mine. In practice, determining that value is difficult. The firm could put the mine on the market for sale, and a buyer would have to determine the value in order to make an offer. But the company could not reasonably expect to put the mine up for sale on a regular basis just to determine its market value. Thus, it would be very difficult to generate a VAR for a copper mine. As an alternative, the firm could use a technique called Cash Flow at Risk, or CAR, (also sometimes known as CFAR). CAR would be defined in terms of an expected cash shortfall. For example, suppose that the copper mine was expected to generate a cash flow of \$10 million per year with a standard deviation of \$2 million a year. The CAR at a probability of 5 percent would be 1.65(\$2 million) = \$3.3 million. This means that

$$\text{Prob}[\$10 \text{ million} - \$3.3 \text{ million} > \text{Actual Cash Flow}] = 0.05.$$

In other words, the probability that \$10 million − \$3.3 million = \$6.7 million will exceed the actual cash flow is 5 percent. That is, the probability that the actual cash flow will be less than \$6.7 million is 5 percent. Note that CAR is calculated in a slightly different manner than VAR. It is expressed in terms of a shortfall from the expected cash flow. Since deviations from the expected value are on average equal to zero, we calculate the CAR without adjusting for the expected cash flow. In other words, the CAR is simply based on the standard deviation times 1.65 (assuming a desired probability level of 5 percent), because the average deviation from the expected value is zero.

Although CAR can be more applicable to certain types of companies, most of the same issues and problems associated with VAR apply equally to CAR.

Another alternative to VAR and CAR is the concept of Earnings at Risk, or EAR. For companies that are concerned about shortfalls in earnings per share, EAR can be used to measure the risk. Of course, we know that value and cash flow are far more important than earnings, because the earnings are a reflection of assumptions about accounting methods, and they do not take risk or the time value of money into account.

Regardless of whether one uses VAR, CAR, or EAR, a good risk manager will collect data over the risky time period and evaluate after the fact whether the measure was a good one. For example, if the VAR is \$1 million for one day at 5 percent, then the risk manager can expect to see a loss of at least \$1 million about one day in a month. Over a period of many months, perhaps even many years, the risk manager can determine whether \$1 million is a reasonable reflection of the true risk.

We have now completed our examination of market risk. We next turn to the other major type of risk: credit risk.

MANAGING CREDIT RISK

Earlier in this chapter, we mentioned a party that could borrow at a fixed rate of 10 1/2 percent or at a floating rate of LIBOR plus 1.1 percent. That party issues floating-rate debt and enters into a swap in which it pays a fixed rate of 9 1/4 percent and receives a floating rate of LIBOR, for a net fixed rate of 10.35 percent, a savings of 15 basis points over issuing fixed rate directly.

If the firm had borrowed at a fixed rate, it would have assumed no credit risk. A lender cannot default. If it borrows at a floating rate and swaps it into a synthetic fixed-rate loan, it faces the risk that the swap dealer will default, leaving it owing the floating rate of LIBOR. Perhaps the 15 basis point saving is simply compensation for bearing the risk that the dealer will default.

The risk of default is faced by any party that may receive obligatory payments from another party. This risk, called credit risk or default risk, is faced by any lender. We have mentioned it many times earlier in the book, especially in conjunction with over-the-counter derivatives. Because banks are used to making loans, their derivatives business is not taking on a different kind of risk than one they already have years of experience with. Since the notional principal on a derivative is not at risk, however, the credit risk of derivatives is typically much lower than the credit risk of loans. Nonbank end users typically do not make loans. When they become users of swaps, however, they become creditors.

Earlier in the book we commented on credit risk whenever we discussed over-the-counter instruments. Futures and exchange-listed options are insured against credit risk by the clearinghouse. Since there has never been a default by the clearinghouse, these contracts can be considered credit-risk free. Over-the-counter contracts are subject to credit risk that varies from party to party, from type of contract to type of contract, and from one point in time to another point in time.

In the bond market, credit risk is typically assessed by examining the credit ratings of issuers. Credit ratings are provided by several firms, most notably Standard & Poor's, Moody's, and Fitch's. Although each agency has its own labels, terms like "Triple-A" and "B-double AA" are used, with the more A's the better. Thus, you would see ratings like AAA, AA, A, BAA, BBB, BB, B, and so on down to D. Analysts for these companies base their ratings on a variety of factors, most notably the financial health of the issuer, and also the states of the economy and the issuer's industry.

To gain a better understanding of what credit risk really involves, let us take a look at the issues underlying credit and default. Interestingly, much of what we know about this subject comes from option pricing theory.

Credit Risk as an Option

Let us consider a very simple firm that has assets with a market value of A_0 and a single zero coupon debt issued with a face value of F. That is, the amount F is due when the bonds mature at time T. The debt has a market value of B_0. Thus, the market value of the stock is

$$S_0 = A_0 - B_0.$$

When the debt matures, the firm will have assets worth A_T and will owe the amount F. If $A_T > F$, the firm will pay off the debt, leaving the amount $A_T - F$ for the stockholders. Thus, the value of their claim at T is $S_T = A_T - F$. If the value of the assets is not adequate to pay off the debt, the creditors will receive the assets worth A_T and the value of the stockholders' claim is zero. In other words, the firm's assets are fully paid out either in the form of partially repaying the creditors or paying the creditors in full, leaving the remainder for the stockholders. Table 15.5 illustrates these results.

Table 15.5 Payoffs to the Suppliers of Capital to the Firm

		Payoffs of Bonds and Stock	
Source of Capital	Market Value at Time 0	$A_T < F$	$A_T \geq F$
Bonds	B_0	A_T	F
Stock	S_0	0	$A_T - F$
Total	$B_0 + S_0$	A_T	A_T

Notice that the payoff to the stockholders is like a call option. In fact, we could write it as follows

$$S_T = \text{Max}(0, A_T - F),$$

which should look like a call option worth S_T at expiration, in which the underlying is an asset worth A_T at expiration, and F is the exercise price. Ignoring the difference between the symbols we use here and those used in Part I, it should be easy to see that stock is indeed a call option and could be priced using the Black-Scholes-Merton model. But before we do that, let us take a further look at the nature of these claims.

If stock is indeed a call option, then we should be able to use put-call parity to relate it to some type of put option. Recall from Chapter 3, that put-call parity for standard options is $P_e(S_0,T,X) = C_e(S_0,T,X) - S_0 + X(1 + r)^{-T}$. In the present situation, we shall use the notation S_0 for the current value of the call, F for the exercise price, and A_0 for the asset price, which is the underlying. Thus, put-call parity would look as follows:

$$P_0 = S_0 - A_0 + F(1 + r)^{-T}.$$

But what exactly is the put? In general, a put is the right to sell the underlying at the exercise price. It is not apparent what the put represents when the stockholders claim is viewed as a call option. Let us rearrange put-call parity:

$$S_0 = A_0 + P_0 - F(1 + r)^{-T}.$$

By definition, the value of the assets is the market value of the stock plus the market value of the bonds: $A_0 = S_0 + B_0$. This means that

$$S_0 = A_0 - B_0.$$

Using the previous two equations, we see that

$$B_0 = F(1 + r)^{-T} - P_0.$$

Thus, the market value of the bonds, which are subject to default, is the market value of a risk-free bond and a short put. The bondholders have, in effect, written an implicit put option to the stockholders. So the stockholders' claim is a claim on the assets, minus a risk-free bond, plus a long put.

The implicit put written by the bondholders to the stockholders is really just a way to specify the limited liability feature that characterizes the nature of corporate ownership. A corporation can borrow money, promising to pay a given amount. If the corporation is unable to pay this amount, it can discharge its obligation in full by turning over (selling) the assets to the bondholders. The personal assets of the owners are not at risk. Thus, their claims are limited to a value of zero. The right held by the stockholders to default without putting their personal assets at risk is regarded in legal circles as limited liability. In economic terms, this right is a put option written by the bondholders to the stockholders.

Because equity is a call option, it can be valued using the Black-Scholes-Merton model. We insert the asset value as the underlying, and the bond face value as the exercise price. The risk-free rate and time to expiration are obvious. The volatility is the volatility of the log return on the assets. The Black-Scholes-Merton model for S_0 is, thus,

$$S_0 = A_0N(d_1) - Fe^{-r_cT}N(d_2).$$

The Black-Scholes-Merton model can also be used to value the risky bond. Recall that $B_0 = A_0 - S_0$. Thus, the market value of the bond can be found as

$$B_0 = A_0 - (A_0N(d_1) - Fe^{-r_cT}N(d_2)) = Fe^{-r_cT}N(d_2) + A_0(1 - N(d_1)).$$

We see that the value of the bondholders' claim is equivalent to the value of a partial claim on a risk-free bond paying F at maturity plus a partial claim on the assets. This should make sense in that there is some probability that the bondholder will receive the full face value and some probability that the bondholder will receive some of the assets, which will be less than the face value.

Let us now turn to the question of how credit risk arises in derivatives.

Credit Risk of Derivatives

There are two types of credit risk in derivatives. Current credit risk is the risk to one party that the other party is unable to make payments that are currently due. Current credit risk is faced by only one party in a derivatives transaction, the party who holds the contract as an asset. In other words, the amount due to that party is positive so it is subject to the possibility that the other party will default. The other party currently owes more than is owed to it, so it faces no current risk of default. Potential credit risk is the risk that the counterparty will default in the future.

Both current and potential credit risk of over-the-counter options are faced by the buyer only. Since the buyer pays the premium to the seller and does not have to do anything else, the seller faces no credit risk. The buyer, however, may eventually choose to exercise the option, by which time the seller may be bankrupt. FRAs and swaps have two-way credit risk. Each party is obligated to do something for the other. Holding other things constant, the credit risk of an FRA or interest rate swap is determined by the credit quality of the counterparty, the terms of the contract, and the shape of the term structure. Consider, for example, a five-year plain vanilla interest rate swap. With an upward-sloping term structure, the implied forward rates are rising. These rates can be equated to the floating rates and will equal the eventual floating rates if there is no shift in interest rates during the life of the swap. When the swap is first established, the party paying the fixed rate and

DERIVATIVES TOOLS

Concepts, Applications, and Extensions
What Derivatives Tell Us About Bonds

Here in our treatment of credit risk, we learned that the stock of a firm that has issued debt is like a call option on the assets of the firm. We saw that a bond subject to credit risk can be viewed as a risk-free bond and a short put option on the assets. In other words, the bondholders have implicitly written a put option to the stockholders. In legal terms, this put option is the principle of limited liability: the claims of the bondholders cannot be met with the personal assets of the stockholders. The bondholders must accept either the amount owed to them or the value of the assets of the firm. Thus, the stockholders can discharge their liability by effectively selling the firm to the bondholders.

As we saw in this chapter, we can determine the market value of the bonds. Let us work an example. Consider a company that has assets worth 1,000. It has one issue of zero coupon bonds outstanding, which matures in two years. The bonds have a face value of 400. The volatility of the assets is 25 percent, and the risk-free rate is 4 percent. Let us use the Black-Scholes-Merton model to determine the market value of the stock, treating it as a call option on the assets. The following shows how this is done, with calculations obtained using the spreadsheet BSMbin7e.xls.

Black-Scholes-Merton variable	Variable in Credit Risk Model	Value
S_0	A_0	1000
X	F	800
r_c	r_c	0.04
σ	σ	0.25
T	T	2
$C_0(S_0,T,X)$	S_0	294.17

Plugging in to the Black-Scholes-Merton model, we would obtain a value as indicated above of 294.17. Thus, the market value of the bonds would be

$$B_0 = A_0 - S_0 = 1{,}000 - 294.17 = 705.83.$$

Let us now determine the yield on the bonds. The bonds have a current value of 705.83 and pay off 800 in two years. The continuously compounded yield, y, can be found using the following relationship:

$$705.83 = 800e^{-y2}.$$

Solving for y gives

$$y = \frac{\ln\left(\frac{800}{705.83}\right)}{2} = 0.0626.$$

Note that this bond offers a risk premium of 2.26 percent over the risk-free rate. This value is typically called the yield spread. The size of the yield spread reflects the credit risk assumed by the lender.

By varying the inputs we can observe how the yield spread changes. In the table below, we see how the yield spread varies with the maturity of the bond, holding everything else constant.

Maturity	Value of Stock	Value of Bonds	Bond Yield (%)	Yield Spread (%)
1	247.79	752.21	6.16	2.16
2	294.17	705.83	6.26	2.26
3	334.29	665.71	6.13	2.13
4	369.98	630.02	5.97	1.97
5	402.34	597.66	5.83	1.83

The relationship between a bond's yield (or yield spread) and maturity is often called the credit risk structure of interest rates. In this case the credit risk structure is increasing out to two years, but then it begins to decrease. This pattern means that over the near horizon, the debt is somewhat riskier. The risk over the longer horizon, however, is smaller, because the firm has more time to improve its financial condition.

The table below shows how the yield spread changes if the amount of the debt increases. In other words, we see below how the yield spread changes if the company has more debt.

Face Value of Debt	Value of Stock	Value of Bonds	Bond Yield (%)	Yield Spread (%)
400	630.91	369.09	4.02	0.02
600	451.23	548.77	4.46	0.46
800	294.17	705.83	6.26	2.26
1000	176.77	823.23	9.73	5.73
1200	100.08	899.92	14.39	10.39

Naturally the more debt the company carries, the larger the yield spread. Note, however, that the firm could carry debt with a face value of \$1,200, which exceeds the value of the assets of \$1,000, and not be bankrupt. The debt would have a value, as indicated above, of about \$900, reflecting the fact that the debt is not due for two more years and that the company could easily be in better financial condition by that time.

In the table below, we examine how the yield spread changes if the volatility of the assets changes.

Volatility	Value of Stock	Value of Bonds	Bond Yield (%)	Yield Spread (%)
0.05	261.51	738.49	4.00	0.00
0.15	267.75	732.25	4.42	0.42
0.25	294.17	705.83	6.26	2.26
0.35	330.66	669.34	8.92	4.92
0.45	370.91	629.09	12.02	8.02

Note that the more volatile the assets, the riskier the debt, as we might expect.

The Black-Scholes-Merton model can be very useful for evaluating credit risk. In fact, one well-known company, the KMV Corporation of San Francisco, has a service that provides credit risk information based on the application of the Black-Scholes-Merton model in the manner shown here.

receiving the floating rate will initially pay more than it receives, but later the floating receipts will begin to overtake the fixed payments. At any given time current credit risk is faced by one party only, which is the party to whom the contract has positive value. Potential credit risk will reflect the combined effects of the risk of changes in interest rates and the risk of the counterparty's assets over the remaining life of the contract. Both parties face potential credit risk.

In the case of interest rate swaps, the potential credit risk is largest during the middle of the swap's life. Potential credit risk is low at the beginning of the swap's life because it is assumed that the swap would not be entered into if there were any significant problems on the near horizon. Potential credit risk is high during the middle part of the swap's life because the counterparty has had time to have its financial condition deteriorate. Potential credit risk is low during the latter part of the swap's life because there are fewer payments remaining.

Recall that interest rate swaps do not involve the exchange of principal but currency swaps often do. This also keeps the credit risk low at the end of the life of the interest rate swap but raises it somewhat for the currency swap.

The explicit prices of derivative transactions do not necessarily differ according to the credit risk of the transaction. Many derivatives dealers and end users will transact only with customers of a minimum acceptable credit quality. Thus, counterparties who barely pass the hurdle pay the same as counterparties who are essentially risk free. Assessment of a counterparty's credit risk is done much the same way as when extending credit. Financial statements and ratings are used along with simulations and other statistical techniques. VAR can even be adapted to accommodate credit risk, though it is much more difficult to do. Stress testing is commonly used to assess credit risk by assuming that extreme economic conditions occur and estimating the likelihood of default.

Although qualifying parties pay the same explicit rate on a derivatives transaction, there are some implicit methods of differentiating among parties with different credit risks.

The primary method used to manage credit risk is to limit the amount of exposure to a given party. The higher the credit quality of the counterparty, the greater the amount of business that can be done with it. This is similar to the principle of diversification. A party spreads out its transactions with numerous end users or dealers. An extreme case of this method is to do no business whatsoever with a counterparty whose credit quality is below a minimum level.

Collateral is commonly used in lending and is becoming increasingly used in derivatives transactions. The two parties might agree that each post collateral by depositing cash or securities with a bank or by providing a letter of credit from another bank. The parties may agree that on any settlement date, the party to whom the derivative value is negative will post collateral equal to a given percentage of the notional principal or the value of the contract. Another type of arrangement is for neither party to post collateral but if either party's credit is downgraded by one of the rating agencies, that party will post a specific amount of collateral. Some transactions are structured so that a ratings downgrade gives the other party the right to terminate the contract at that point, meaning that the value of the swap is due and payable at that time.

We discussed marking to market extensively when we studied futures markets. Marking to market is designed to reduce credit risk by forcing losing parties to pay winning parties before losses accumulate to a large amount. Some derivative contracts are structured to have marking to market on a periodic basis. For example, a swap that settles semiannually might stipulate that the contracts are marked to market every two months. When the contract is marked to market, the market value is calculated. The party to whom the value is negative owes that full amount in cash to the other party. The fixed rate is then reset to the current market fixed rate. The actual settlement, meaning the exchange of interest payments, proceeds at the regularly scheduled dates though with a new fixed rate. This procedure forces the party on the losing side to pay up before the accumulated loss is potentially much greater.

One should not always assume that the party with the weaker credit is the dealer. Many dealers are banks whose credit ratings are held down by the possibility of loan losses in their domestic and global banking business. In response to the concerns of some end users, a number of dealers have set up separate subsidiaries who do exclusively swap and derivatives transactions with end users. These subsidiaries are provided a large capital cushion and are not responsible for the debts of the parent company. Consequently, the subsidiary would have a higher credit rating than the parent company and most of these subsidiaries are rated AAA. These types of firms are called enhanced derivatives products companies and sometimes special purpose vehicles (also known as special purpose entities). When these firms were first created several years ago, they were greeted with much enthusiasm. The amount of business they have generated, however, has been surprisingly small. In fact, many parent companies also maintain derivative groups within the parent company as well as the subsidiary. Most end users were content to do business with the parent company as long as its rating was reasonably high, though not necessarily AAA.

Another procedure widely used to reduce default risk is netting. Because it is used so much and it is so important, we shall devote a special section to it.

Netting

Netting is a term to describe several similar processes that are used to reduce the amount of cash paid from one party to the other by deducting any cash owed by the latter party to the former party. Netting between two parties is called bilateral netting while netting across several parties is called multilateral netting. Multilateral netting is effectively what happens when there is a clearinghouse, such as the options or futures clearinghouses we discussed in Chapters 2 and 8.

There are several forms of netting. We have already seen payment netting, in which two parties who owe each other money on a given day agree to simply determine the net amount owed and have one party pay the other. As you should remember, this is a fairly standard procedure in interest rate swaps. If party A owes \$1,250,000 to party B who owes \$890,000 to party A, then party A simply makes a payment of \$360,000 to party B. This eliminates the possibility of party B defaulting, at least on this payment.

In cross product netting, payments for one type of transaction are netted against payments for another type of transaction. Suppose, for example, that on the date on which party A owes \$260,000 to party B, party B also owes A \$175,000 as the interest on a loan. If the two parties had agreed to cross product netting, then A would owe B only \$260,000 − \$175,000 = \$85,000. Of course, the two parties might seldom have transactions with payments due on the same day. Consequently, cross product netting usually comes into play in bankruptcy or in a takeover involving two parties, each owing the other various amounts arising from different transactions.

If the parties agree to netting by novation, the net value of their mutual obligations is replaced by a single new transaction. In our example above, if A owes B \$85,000, then the transactions that have an aggregate value of \$85,000 to B are replaced by a new transaction with a value of \$85,000. Netting by novation is used more often in the foreign exchange markets.

Closeout netting is the stipulation that if default occurs, only the net amount is owed. This greatly reduces the credit risk by reducing the amount of cash owed by a defaulting party by the amount owed to it. Consider the following example:

Company XYZ and dealer FinSwaps are engaged in four transactions with each other. From the dealer's perspective, the market values are as follows:

Swap 1	−$1,179,580
Swap 2	+$1,055,662
Option 1	+$1,495,255
FRA 1	−$ 892,530
	+$ 478,807

XYZ defaults with no closeout netting:

Suppose XYZ demands payments on swap 1 and FRA 1 and refuses to pay on swap 2 and option 1. FinSwaps owes XYZ $2,072,110, the sum of the values of swap 1 and FRA 1. This process is commonly called cherry picking because the bankrupt party selects the attractive transactions and walks away from the others.

XYZ defaults but closeout netting is used:

XYZ owes FinSwaps $478,807.

Today, there is hardly an issue over the use of closeout netting. It is standard in virtually every swap in most countries, and most countries' bankruptcy laws have been revised to recognize the legitimacy of closeout netting.

In the presence of netting, the credit risk of swaps is clearly limited to the net payments, which is less than the present value of the interest payments on a loan of equivalent notional principal. Because of this and the fact that the over-the-counter derivatives market has imposed strict requirements to reduce the credit risk, its record in the matter of defaults is excellent. The default of the London borough Hammersmith and Fulham in 1988 was the result of a legal ruling that permitted it to walk away from losing transactions and accounted for about half of all defaults in the over-the-counter market until that time. The record since that time has been quite good as well. In fact, the derivatives market has established a much better record on the matter of default than has the commercial loan market. Clearly the concerns expressed by some over the large size of the market (as exemplified by the notional principal outstanding) overstate the real risk.

Credit Derivatives

In the early 1990s a major innovation in credit risk management appeared on the scene: the credit derivative. A credit derivative is an instrument designed to separate market risk from credit risk and to allow the separate trading of the latter. Specifically it is a derivative or derivative-like instrument with a payoff determined by whether a third party makes a promised payment on a debt obligation. Credit derivatives permit a more efficient allocation and pricing of credit risk. Parties who wish to rid themselves of credit risk can engage in credit derivative transactions to pass the credit risk on to another party who is willing to accept it. Ultimately, this transferal of risk greatly benefits those who borrow and lend, as well as those who transact in derivatives that are subject to default, for it helps assure them that the premiums associated with the risk of default are appropriate for that level of risk. Banks, in particular, use credit derivatives both as buyers and sellers. Banks must limit the amount of exposure to a particular borrower, to a particular geographic area, or to a particular industry. They can often provide a loan, thereby servicing the borrower, but can then sell the credit risk using a credit derivative.

Of course, for any party that eliminates credit risk, there must be a party that accepts it. Purchasing credit risk that is unrelated to the risks already in a portfolio can help diversify the portfolio. In addition, speculators could be willing to accept the credit risk, believing that it is less than what is priced into the contract. The primary parties who accept the credit risk transferred through credit derivatives are other banks, insurance companies, and hedge funds.

A credit derivative transaction involves three parties. One is the credit derivatives buyer, which is a party that holds credit risk that it wishes to eliminate. The second party is the credit derivatives seller, which is the party that is willing to acquire the credit risk. The third party is the party on whose credit the transaction is based. This party is called the reference entity. For example, suppose Bank A has made a loan to Company C.

Now Bank A has too much exposure to the credit of Company C and wishes to sell the credit risk. It finds another company willing to take on the risk, which happens to be a hedge fund called Fund B. Bank A is the credit derivative buyer, Fund B is the credit derivative seller, and Company C is the reference entity. The reference entity is involved in the credit derivative in name only. In fact, it may well be unaware that it is the reference entity in a credit derivatives transaction. It is the right of the party exposed to the credit risk of the reference entity to sell that risk to another party without involvement of the reference entity.

Credit derivatives are somewhat more complex than other derivatives for a variety of reasons. Credit risk is more difficult to measure, and losses from credit events are less frequent than losses from movements in, say, interest rates or exchange rates. Hence, it is more difficult to build historical models of credit losses. Credit derivatives are also less liquid instruments than ordinary derivatives, and they involve complex legal issues related to contract terms and documentation as well as to definitions of what constitutes a credit loss.

For example, there are numerous types of events that might be specified as a credit loss in the contract. These include not making an interest or scheduled principal payment following a reasonable grace period, declaration of bankruptcy, repudiation of a liability, restructuring of the borrower's other debts, or acceleration of payment of another liability. The latter could be caused by a provision in another liability that permits another creditor to require premature payment of a debt. These possible events must be spelled out in the contract.

We should also remember that a credit derivative is itself subject to credit risk. If the party promising to pay in the event of a credit loss is unable to pay, there is a credit loss on a credit derivative. The potential for a default by both the reference entity and the seller is, however, usually very small.

Nonetheless, the credit derivatives market is large and growing rapidly. The Bank for International Settlements surveys that we have previously cited have only recently begun documenting the size of the credit derivatives market. Its surveys, however, are limited to only one type of credit derivative (credit default swaps, which we will cover later). As of December 2005, the notional principal of these instruments is estimated at \$10.3 trillion, but keep in mind the misleading nature of notional principal as a measure of risk. In the case of credit derivatives, this figure represents the aggregate face value of bonds and loans covered by credit derivatives. The market value of these credit derivatives is estimated at \$346 billion. Some other information is available from the United States Office of the Comptroller of the Currency, which regulates federally chartered banks in the United States. Its surveys, obtainable from its Web site, http://www.occ.treas.gov, indicate that the notional principal of credit derivatives at these large U.S. banks was \$5.8 trillion at the end of December 2005.

Now let us take a look at the various types of credit derivatives. In the discussions that follow, we will use the term *buyer* to refer to the credit derivatives buyer who is the party who wants to eliminate the credit risk, and *seller* to refer to the credit derivatives seller who is the party willing to accept the credit risk.

Total Return Swaps Perhaps the simplest credit derivative is the total return swap, which is a swap transaction in which the buyer agrees to pay the total return on a particular reference asset, such as a specific bond issued by the reference entity, to the seller. The seller, in turn, agrees to pay the buyer a rate such as LIBOR plus a spread. As in any swap, these payments are made on a regular schedule, usually quarterly or semiannually. The total return includes any interest payments as well as unrealized capital gains. If the bond incurs a capital loss, the buyer then receives a payment from the seller. The buyer is, in effect, promising to pay the seller the bond's total return and to correspondingly receive an interest payment, thereby ridding itself of the risk of price changes caused by factors unrelated to interest rate movements. Presumably such price changes would occur primarily from changes in the reference entity's credit risk. The structure of a total return swap is illustrated in Figure 15.4. Note that unrealized capital gains and losses are non-cash amounts, but the swap requires cash payment for unrealized capital gains and losses.

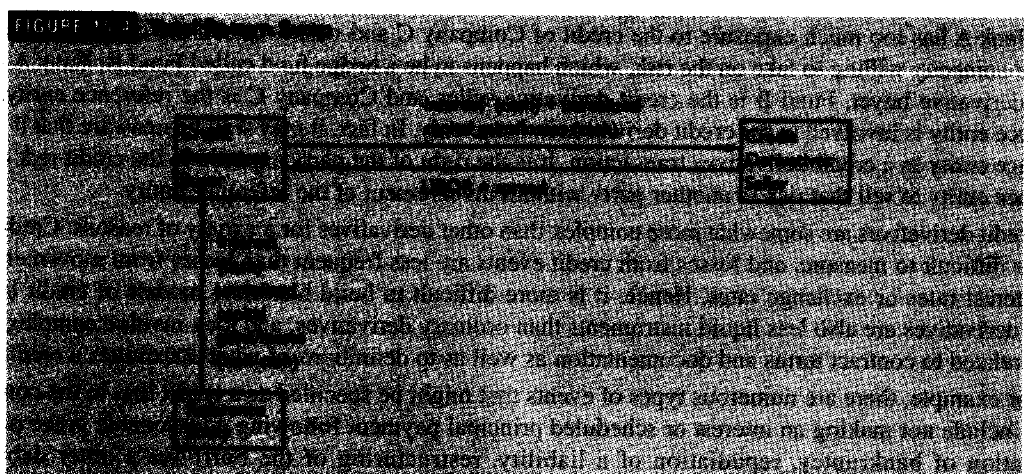

The total return swap maintains the bond in the possession of the buyer. The seller is seeking the return on the bond without buying the bond. Note that market risk due to interest rate movements is not completely eliminated since those movements are reflected in the LIBOR payment. Incidentally, the LIBOR payment can, alternatively, be a fixed rate or it can be converted from LIBOR to a fixed rate with a plain vanilla swap. Note that the total return swap is much like an ordinary swap, but the interest payment corresponds to that on a specific bond and the unrealized capital gains and losses are paid.

Total return swaps are practical only for reasonably liquid bonds and not for ordinary loans. This is because the parties must be able to identify the total return over the settlement period in order to determine how much the buyer owes the seller. Consequently, there must be a reliable measure of the market value of the bond, which can occur only if there is an actively traded market for the bond or a dealer willing to quote a price for the bond.

Credit Default Swaps A second type of derivative is the credit default swap, sometimes called a credit swap. The buyer holds a bond or loan issued by the reference entity and enters into the credit default swap, which obligates it to make a series of periodic payments to the seller. If a credit event occurs to the reference entity, the seller compensates the buyer for the loss. This compensation can take one of two primary forms. In a *cash settlement*, the seller pays the buyer the difference between the amount defaulted less any amount recovered. In *physical delivery*, the buyer simply delivers the reference instrument to the seller. Of course, in earlier chapters we discussed the use of cash settlement and physical delivery for other derivative contracts. One other type of settlement, called *fixed-settlement*, is also used, though infrequently. In a fixed settlement, the seller pays the buyer a fixed amount regardless of the amount of the loss.

Although a credit default swap is called a swap, it is easy to see why it is essentially an insurance policy.[8] The buyer pays the seller a periodic fee, which is similar to an insurance premium. The seller then compensates the buyer in the event that a credit loss occurs. Credit swaps can be written on a reference entity's bonds or loans. A credit default swap is illustrated in Figure 15.5.

Credit default swaps are probably the most widely used credit derivative, and there are many variants. One of the primary differentiating features is how the reference entity is specified. The type of credit default swap we have been describing is on a single reference entity. There are other credit default swaps in which

[8]In the United States, insurance is regulated at the state level. Hence, any product called *insurance* requires regulatory approval from all fifty states before it can be generally offered. Credit derivatives serve much the same purpose as credit insurance but are private contracts and are not directly regulated by the federal or state governments.

FIGURE 15.5 Credit Default Swap

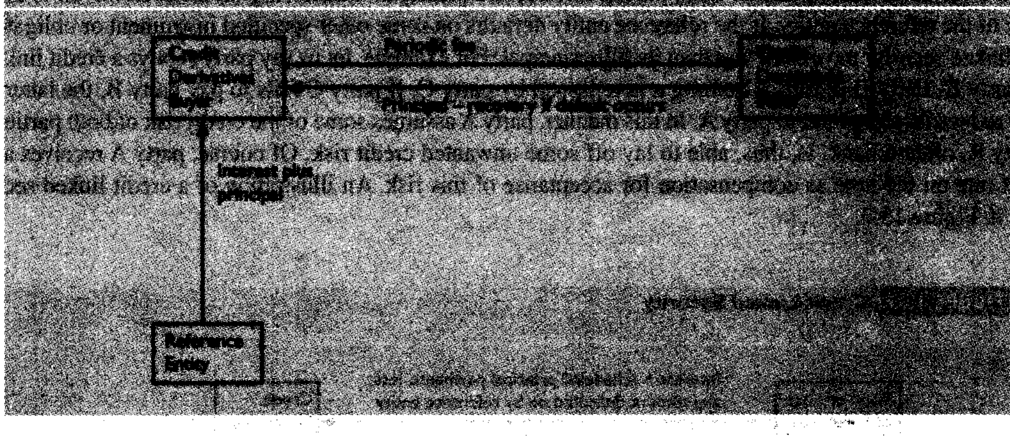

the reference entity is actually a portfolio of loans or bonds of various reference entities. This portfolio is usually referred to as a basket. For a credit default swap on a basket, the swap is specified to pay off when the n^{th} party-to-default actually defaults, where n is a contractually specified number. That is, a first-to-default credit default swap will pay off when the first entity in the basket defaults. A second-to-default credit default swap does not pay off until the second default occurs. The payoff is based only on the second default and does not cover losses from the first default. Therefore, n can be any number up to the total number of reference entities in the basket.

Credit Spread Options A credit spread option is just an option on the spread of a bond over a bond issued by the reference entity. Normally in the bond market, the spread of a bond's yield over the yield on a comparable maturity U.S. Treasury instrument is called the yield spread. It fluctuates with investors' perceptions of the credit risk in the market. Suppose the credit risk of the reference entity currently dictates that it should pay 50 basis points in yield over a U.S. Treasury note of comparable maturity. The buyer might buy a credit spread option with an exercise

FIGURE 15.6 Credit Spread Option

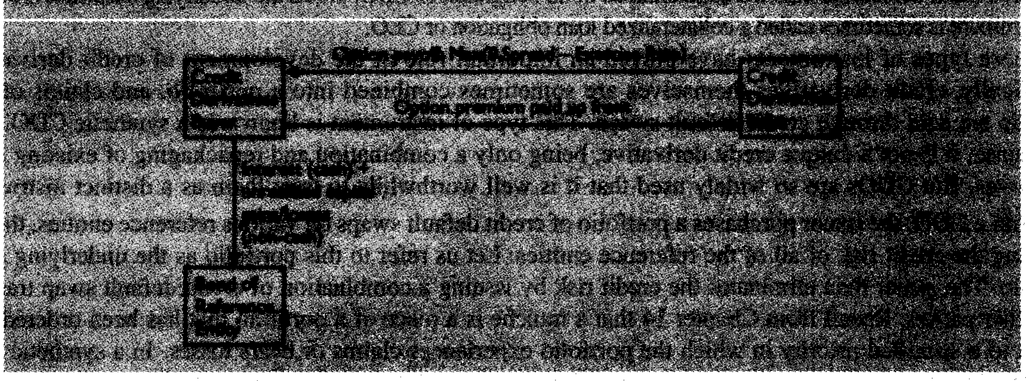

rate of 60 basis points. At expiration, or early if the option is American style, the option is in-the-money if the credit spread exceeds 60 basis points. Of course, this will occur only if the market perceives that the credit risk has increased. For this right, the buyer pays the seller a premium up front. Thus, this instrument is like an ordinary option, but the underlying is the spread of the bond over the otherwise comparable U.S. Treasury security. This type of credit derivative also requires that the underlying bond be sufficiently liquid such that a reliable assessment of its credit spread can be obtained. An illustration of a credit spread option is shown in Figure 15.6

Credit Linked Security Another type of credit derivative is the credit linked security. This type of instrument appears to be an ordinary bond or note, paying interest and principal. Underlying the note, however, is the credit quality of the reference entity. If the reference entity defaults on some other specified instrument or obligation, the credit linked security pays back less than its full principal. For example, let us say party A buys a credit linked note from party B. Party B in turn is holding a note issued by party C. If party C fails to pay party B, the latter is then able to reduce its obligation to party A. In this manner, party A assumes some of the credit risk of both parties A and B. Party B, often a bank, is, thus, able to lay off some unwanted credit risk. Of course, party A receives a higher interest rate on the note as compensation for acceptance of this risk. An illustration of a credit linked security is shown in Figure 15.7.

FIGURE 15.7 Credit Linked Security

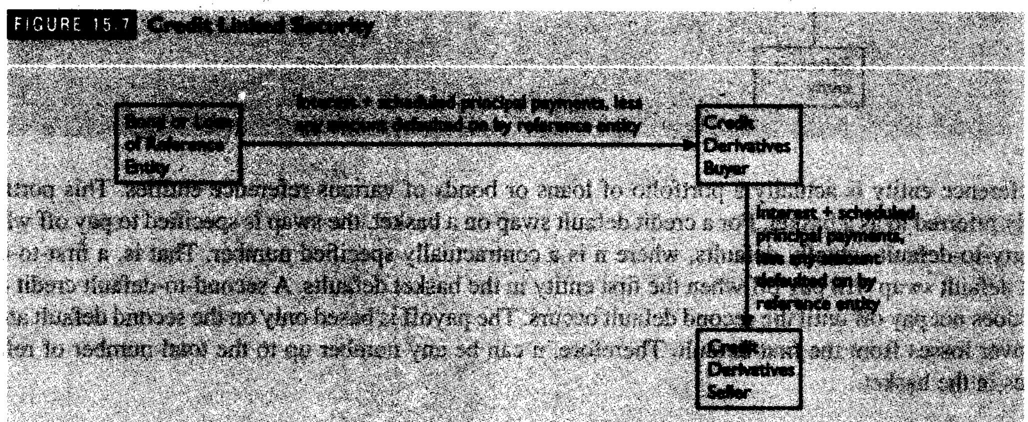

Synthetic CDOs In Chapter 14, we discussed asset-backed securities. We illustrated how mortgages are combined into portfolios, and how claims on the interest and principal from the mortgages are sold in the form of pass-throughs and collateralized mortgage obligations, or CMOs. Recall that in a CMO, claims on the underlying portfolio are divided into prioritized securities called tranches. When the underlying securities are bonds or loans instead of mortgages, this instrument is called a collateralized debt obligation or CDO. When the underlying securities are bonds, the instrument is sometimes called a collateralized bond obligation or CBO. When the underlying securities are loans, the instrument is sometimes called a collateralized loan obligation or CLO.

These types of instruments have played an important role in the development of credit derivatives. Specifically, credit derivatives themselves are sometimes combined into a portfolio, and claims on that portfolio are sold through credit default swaps. This type of instrument is known as a synthetic CDO.[9] In a strict sense, it is not a unique credit derivative, being only a combination and repackaging of existing credit derivatives. But CDOs are so widely used that it is well worthwhile to treat them as a distinct instrument.

With a CDO, the issuer purchases a portfolio of credit default swaps on various reference entities, thereby assuming the credit risk of all of the reference entities. Let us refer to this portfolio as the underlying credit portfolio. The issuer then eliminates the credit risk by issuing a combination of credit default swap tranches with other parties. Recall from Chapter 14 that a tranche is a piece of a portfolio that has been ordered with respect to a specified priority in which the portfolio experiences claims or bears losses. In a synthetic CDO these tranches assume, in a specified order of priority, the credit losses from the underlying credit portfolio. The highest priority tranche bears credit losses last, with each tranche below it bearing credit losses

[9]CDOs wherein the underlying is a portfolio of securities are now sometimes called cash CDOs to distinguish them from synthetic CDOs, wherein the underlying is a portfolio of credit derivatives.

FIGURE 15.8 Synthetic CDO

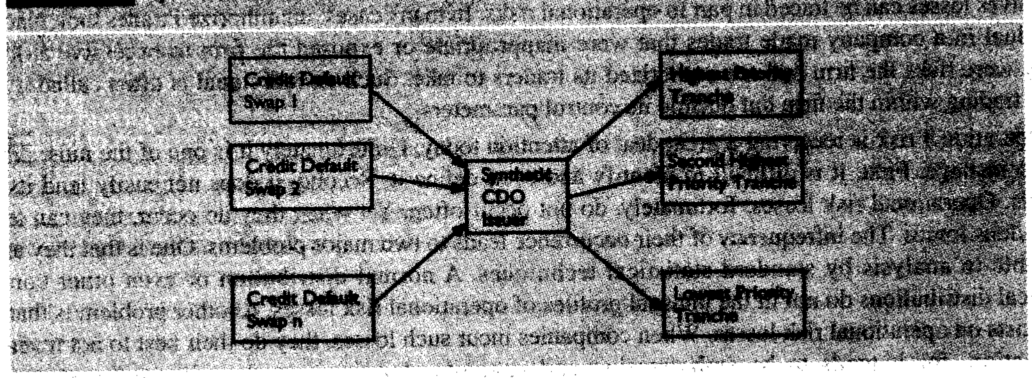

successively earlier. At least one tranche will probably bear little credit risk because all of the other tranches will assume any losses ahead of it. The lowest priority tranche will bear the most credit risk, because it will have to assume the first credit losses from the underlying credit portfolio. Other tranches lying in between the two extremes bear various degrees of credit risk. Figure 15.8 illustrates a synthetic CDO.

The credit risk assumed is, of course, a major determinant of the pricing of these tranches. Parties purchasing the riskier tranches pay less and have higher expected returns because of the risk. The pricing of credit derivatives is quite complex. The basic idea stems from the notion of default as an option, as discussed earlier, but while the theory is sound, implementing it is difficult. Moreover, the pricing of credit derivatives is complicated by the fact that the payoffs are often contingent on the credit risk of multiple entities. Hence, the correlations between the respective credit risks of the various entities are important and quite difficult to evaluate.

In spite of the complexities of credit derivatives, the market is, as noted earlier, growing rapidly. In fact, credit derivatives have become such a popular instrument that a number of organizations have created indices of the performance of credit derivatives. These indices are useful measures of the general level of credit risk in the economy. In addition, it is possible to buy and sell the indices as one would an ordinary security. Thus, in addition to engaging in direct transactions in credit derivatives, market participants can buy and sell credit risk as though it were a mutual fund.

OTHER TYPES OF RISKS

Although market risk and credit risk comprise the two primary risks, they are by no means exhaustive. The unforeseen circumstances that can occur touch on almost all aspects of running a business.

Operational risk is the risk of a breakdown in the operations of the derivatives program or risk management system. This could include such events as power failures, computer problems such as viruses and software bugs, the failure of staff personnel to monitor and record transactions properly, the failure of staff personnel to have the necessary knowledge of potentially complex transactions, the failure to have proper documentation, and fraud perpetrated by traders or staff personnel. These risks are present in virtually any type of business operation, but because derivatives transactions are generally somewhat complex and usually involve large amounts of money, the need to avoid these problems is critical.

Perhaps the most important part of operational risk is the effectiveness of proper controls. For example, if certain persons are authorized to engage in derivatives transactions, then it is imperative that their activities be monitored by personnel who report to persons higher in the organization. Anyone authorized to trade must be restricted in what contracts and how much that person can trade. Controls must be in place to make sure that these policies are followed. As obvious as that seems, some firms have not done a good job of ensuring that well-established policies and procedures are actually followed. In fact, it is safe to say that many derivatives losses can be traced in part to operational risks. In many cases, unauthorized trades took place. An individual in a company made trades that were inappropriate or exposed the firm to excessive risks that clearly were risks the firm had not authorized its traders to take. Such an individual is often called a rogue trader, trading within the firm but outside its control parameters.

Operational risk is receiving a great deal of attention today. Unfortunately it is one of the most difficult risks to manage. First, it is difficult to identify and even define it. Second, it does not easily lend itself to analysis. Operational risk losses, fortunately, do not occur often. Yet when they do occur, they can lead to tremendous losses. The infrequency of their occurrence leads to two major problems. One is that they are not amenable to analysis by standard statistical techniques. A normal distribution or even other common statistical distributions do not fit the unusual profiles of operational risk losses. Another problem is that little data exists on operational risk losses. When companies incur such losses, they do their best to not reveal this information, for it tends to be embarrassing and suggests sloppiness in their internal systems. The infrequency of past problems and lack of data make it difficult to predict when and where such losses will occur. The risk management industry, however, is beginning to recognize this problem and some efforts to share data among firms have begun.

Operational risks are more along the lines of risks that are addressed by traditional insurance policies. When an insurance firm provides protection against a singer's voice during a concert tour, it is insuring against an infrequent event of an operational risk nature. When an insurer provides protection against employee fraud, it is insuring against operational risk. It is not surprising, therefore, that insurance companies have begun to look at marketing products for insuring against operational risk in the financial industry and, in particular, in connection with derivatives-related transactions. While there are no derivative contracts that protect against operational risk, some discussion has begun toward creating operational risk derivatives, and they will surely exist in the near future.

Model risk is the risk that in pricing a financial instrument, such as a derivative, the firm will use an inappropriate model or a model containing errors (including programming bugs and errors), or use incorrect inputs. This is a critical issue for dealers, because offering derivatives at attractive prices and then managing the risk is their business. It can also be critical for end users, because they are less likely to have the personnel and knowledge to properly price the instruments.

We have learned enough in this book to take a quick look at model risk. Suppose a customer wants to buy a three-year American put option on the S&P 500. Let the index be at 1,450, the exercise price be at 1,450, the risk-free rate be 6 percent, the volatility be 18 percent, the dividend yield be 1.5 percent, and, of course, the time to expiration be three years. Because this is an American option, we should use a binomial tree to properly reflect the early exercise feature. It is tempting, nonetheless, to use the Black-Scholes-Merton model, so that is what we mistakenly do. Using your spreadsheet BSMbin7e.xls, we obtain a value for the call of 88.05. Let us get the correct value, however, using the binomial model. Using 200 time steps, we obtain 111.82. This is an error of over 20 percent. If the customer wanted $1,000,000 of notional principal of this option, we would be losing over $200,000 just in selling the option. Further losses would undoubtedly be incurred in trying to manage the risk by delta, gamma, and vega hedging.

Here is another variation of model risk. Suppose we used the correct model, sold the option at 111.82, and then hedged away the risk with other transactions. Now suppose one year later we needed to assign a value to the option. With two years to go and the same S&P value, the put would be worth 99.56. But suppose we used a volatility of 22 percent. The put would have a value of 122.68. This is a significant difference. Yet volatility variations from, for example, 18 percent to 22 percent are common. One might be tempted not to view this as a form of model risk, but any time the inputs are unknown, model risk exists. With an input as critical as the volatility is to an option, the model risk can be tremendous.

Some major financial institutions have discovered embarrassing errors, even in simple valuation problems such as forward contracts, that have led to significant losses. The best insurance against model risk is knowledge: knowledge of the theories and models.

Liquidity risk is the risk that a firm will need to enter into a derivatives transaction and find that the market for that transaction is so thin that the price includes a significant discount or premium for that liquidity. Anyone making a market in an illiquid instrument will bear significant risk and would charge wide bid-ask spreads or even be unwilling to trade. Most plain vanilla derivatives have little such risk but exotic transactions can have significant liquidity risk. Liquidity risk is hardly confined to derivatives, however, as was evidenced in the August–September 1998 period when the famous hedge fund Long-Term Capital Management (LTCM) found itself holding positions in very illiquid bonds from around the world. LTCM had sophisticated risk management systems and some of the smartest people in the business, but it misestimated the effects of illiquidity during a global financial crisis where markets around the world were falling in unison. LTCM suffered severe losses that led to a bailout by a consortium of banks that was orchestrated by the U.S. Federal Reserve.

A great deal of research is going on to model liquidity risk or, at the minimum, to think about its effects when managing other types of risks. Some discussion has begun on the creation of liquidity risk derivatives, which would allow parties to buy and sell liquidity risk just like any other form of risk.

Accounting risk is the uncertainty over the proper accounting treatment of a derivatives transaction. Accounting for derivatives has been a significant source of controversy and risk for many years. Users of derivatives have lived with an ongoing concern that the manner in which they account for derivatives will be declared inappropriate after the fact and that they will be required to restate certain transactions with the potential to lower past earnings figures. The significance, timeliness, and critical importance of proper accounting for derivatives is a reason why we examine this subject in more detail in the next chapter.

Legal risk is the risk that the legal system will fail to enforce a contract. For example, suppose a dealer enters into a swap with a counterparty who, upon incurring a loss, then refuses to pay the dealer, arguing that the dealer misled it or that the counterparty had no legal authority to enter into the swap. These arguments have been successfully used in the legal system, particularly by local governments. The risk of contracts not being enforced is a serious one for dealers.[10] This risk can effectively turn a swap or an FRA into an option, because the counterparty simply walks away without paying if the market moves against it. Not too surprisingly, no counterparty has ever claimed that it was misled or had no authority to do derivatives after making money on them. Efforts to control legal risk largely arise from having good documentation of all transactions. The industry, through its trade association ISDA (International Swaps and Derivatives Association), has established standards of documentation for various derivative transactions, including contract templates, formal definitions of key terms, and specific provisions that are widely used in over-the-counter derivatives transactions. ISDA has also been instrumental lobbying and seeking legal opinions before issues become tied up in litigation.

Tax risk is the risk that taxes or the interpretation of tax laws will change unexpectedly. A good example of this was described in Appendix 11 where the tax status of hedging, which had been established almost a half century prior, was called into question. For a period of time, the possibility existed that certain hedging transactions would be taxed in a different manner, rendering them much less attractive. Moreover, the threat that completed transactions will have to be re-taxed always looms. Tax risk has managed to completely eliminate the use of certain transactions, such as the American Stock Exchange's PRIMES and SCORES, which were precursors to the equity-linked debt we covered in Chapter 14. As another example of tax risk, the U.S. Congress has been considering for many years levying a tax on each futures transaction to pay the cost of regulating the futures markets.

Regulatory risk is the risk that regulations or regulatory philosophy will change. Because regulators are appointed by the political party in control, it is easy for a regulatory agency to go from a free-market, light-regulation approach to a more direct and heavy-headed approach. Regulatory risk means that certain existing or contemplated transactions can become illegal or regulated.

---

[10]The aforementioned Hammersmith and Fulham case illustrates the dangers of legal risk to the dealer. The case was supported on the doctrine of *ultra vires,* which means to have no legal authority.

Settlement risk is commonly faced in international transactions. Let us say that a bank in country A engages in a financial transaction with a corporation in country B in which both the bank and the corporation will be required to pay each other rather than net their payments.[11] The bank and the corporation are twelve hours apart with the bank's day beginning first. On settlement day the bank wires its funds to the corporation under the assumption that when the markets open the corporation will wire its funds to the bank. When the corporation's markets open, however, the corporation announces that it is bankrupt and will suspend all payments. Thus, the bank will be out the money and will have to get in line with the corporation's other creditors. This risk is sometimes called Herstatt risk, named for the German bank that in 1974 failed under similar circumstances. Settlement risk can arise out of bankruptcy, insolvency, or fraud.

This enumeration of risks does not necessarily cover every possible contingency. Foreign currency transactions, for example, are subject to political risk, meaning that the government of a country may take over the financial system and declare certain transactions null and void or refuse to allow currency to leave the country. The Chicago Board of Trade was forced to suspend operations one day due to a flood in downtown Chicago. Risk managers needing to trade CBOT products had to look for substitutes in other markets. It would be impossible to have every contingency covered, but risk managers need to be aware of as many problems as possible that could occur.

There is one final source of risk that is of considerable concern to many people, especially regulators. There is a belief that when one company defaults, it could trigger the default of one of its creditors, which could trigger further defaults. These effects could ripple through the entire financial system, leading to widespread panic and a meltdown of the whole system. This is called systemic risk. In modern times, the closest the world's economy has come to such a breakdown was during the stock market crash of October 1987. On a lesser scale, another example is the global financial crisis that occurred in September 1998. As it turned out, while there was a great deal of panic and enormous losses, the markets did not even come close to a total breakdown. Regulators worry, however, that the volume of derivatives transactions, particularly over-the-counter derivatives, is much greater now. While this might be a legitimate concern in an economy in which most financial institutions do not practice risk management, there is little cause for concern today. Nearly every major financial institution in the world practices risk management and most maintain a relatively hedged position in derivatives. There have been numerous shocks to the financial system in the 1990s but none has come remotely close to causing a systemic breakdown.

Questions and Problems

1. How is the practice of risk management similar to hedging and how is it different?
2. Identify why risk management can be beneficial to stockholders.
3. Explain the difference between market risk and credit risk. Are techniques for managing market risk appropriate for managing credit risk?
4. Identify the three parties involved in any credit derivatives transaction and describe how they differ in their roles and responsibilities with regard to the transaction.
5. If a portfolio of derivatives is delta hedged by adding a position in Eurodollar futures, what other forms of market risk might remain? How can these risks be eliminated?
6. Interpret the following statements about Value at Risk so that they would be easily understood by a nontechnical corporate executive:
   a. VAR of $1.5 million, one week, probability = 0.01
   b. VAR of $3.75 million, one year, probability = 0.05
7. Critique each of the three methods of calculating Value at Risk, giving one advantage and one disadvantage of each.

[11]For example, a currency swap is usually done without netting because the payments are occurring in two different currencies.

8. Comment on the current credit risk assumed for each of the following positions. Treat them separately; that is, not combined with any other instruments.
   a. You are short an out-of-the-money interest rate call option.
   b. You entered into a pay fixed-receive floating interest rate swap a year ago. Since that time, interest rates have increased.
   c. You are long an in-the-money currency put option.
   d. You are long a forward contract. During the life of the contract the price of the underlying asset has decreased below the contract price.
9. Explain how closeout netting reduces the credit risk for two firms engaged in several derivatives contracts.
10. How does the legal system impose risk on a derivatives dealer?
11. Consider a firm that has assets that generate cash but which cannot be easily valued on a regular basis. What are the difficulties faced by this firm when using VAR and what alternatives would it have?
12. How is liquidity a source of risk?
13. Explain how the stockholders of a company hold an implicit put option written by the creditors.
14. Identify the five types of credit derivatives and briefly describe how each works.
15. Suppose your firm is a derivatives dealer and has recently created a new product. In addition to market and credit risk, what additional risks does it face that are associated more with new products?
16. Consider a portfolio consisting of $10 million invested in the S&P 500, and $7.5 million invested in U.S. Treasury bonds. The S&P 500 has an expected return of 14 percent and a standard deviation of 16 percent. The Treasury bonds have an expected return of 9 percent and a standard deviation of 8 percent. The correlation between the S&P 500 and the bonds is 0.35. All figures are stated on an annual basis.
   a. Find the VAR for one year at a probability of 0.05. Identify and use the most appropriate method given the information you have.
   b. Using the information you obtained in part *a*, find the VAR for one day.
17. Calculate the VAR for the following situations:
   a. Use the analytical method and determine the VAR at a probability of 0.05 for a portfolio in which the standard deviation of annual returns is $2.5 million. Assume an expected return of $0.0.
   b. Use the historical method and the following information for the last 120 days of returns to calculate an approximate VAR for a portfolio of $20 million using a probability of 0.05:

Less than −0%	5
−10% to −5%	18
−5% to 0%	42
0% to 5%	36
5% to 10%	15
Greater than 10%	4

18. The following table lists three financial instruments and their deltas, gammas, and vegas for each $1 million notional principal under the assumption of a long position. (Long in a swap or FRA means to pay fixed and receive floating.) Assume that you hold a $12 million notional principal long position in the three-year call option, an $8 million notional principal short position in the three-year swap, and an $11 million notional principal long position in the FRA. Each derivative is based on the 90-day LIBOR.

Instrument	Delta	Gamma	Vega
3-year call option with exercise rate of 0.12	$40	$1,343	$5.02
3-year swap with fixed rate of 0.1125	$152	−$678	$0
2-year FRA with fixed rate of 0.11	$72	−$390	$0

a. As described above, you have three instruments currently in your portfolio. Determine your current portfolio delta, gamma, and vega. Describe in words the risk properties of your portfolio based on your calculations.
b. Assume that you have to maintain your current position in the call option but are free to increase or decrease your positions in the swap and FRA and you can add a position in a one-year call with a delta of \$62, a gamma of \$2,680, and a vega of \$2.41. Find the combination of notional principals that would make your overall position be delta hedged, gamma hedged, and vega hedged.

19. Suppose you own 50,000 shares of stock valued at \$35.50 per share. You are interested in protecting it with a put that would have a delta of –0.62. Assume, however, that the put is not available or is unfairly priced. Illustrate how to construct a dynamic hedge using a risk-free debt instrument that would replicate the position of having the put. Ignore the cost of the puts. Show how the hedge works by explaining what happens if the stock falls by one dollar.
20. Company CPN and dealer SwapFin are engaged in three transactions with each other. From SwapFin's perspective, the market values are as follows:

Swap 1	–\$2,000,000
Forward 1	+\$1,500,000
Option 1	–\$ 500,000
	–\$1,000,000

Explain the consequences to SwapFin if CPN defaults with and without closeout netting. Within your answer, explain what is meant by cherry picking.

21. (Concept Problem) Suppose you enter into a bet with someone in which you pay \$5 up front and are allowed to throw a pair of dice. You receive a payoff equal to the total in dollars of the numbers on the two dice. In other words, if you roll a 1 and a 2, your payoff is \$3 and your profit is \$3 – \$5 = –\$2. Determine the probability associated with a Value at Risk of \$0.
22. (Concept Problem) A company has assets with a market value of \$100. It has one outstanding bond issue, a zero coupon bond maturing in two years with a face value of \$75. The risk-free rate is 5 percent. The volatility of the asset is 0.30. Determine the market value of the equity and the continuously compounded yield on the bond. (Use the spreadsheet BSMbin7e.xls for calculations.)

16

MANAGING RISK IN AN ORGANIZATION

In the first 15 chapters, we studied the use of many types of derivatives to manage risk in a variety of situations. We learned what types of contracts are available and the situations for which each is most appropriate. We also learned that derivatives are powerful instruments that have high degrees of leverage. Moreover, some of these instruments are quite complex. Because of these attendant factors, there is another important component of the process of using derivatives for managing risk: effective risk management requires an effective organizational structure. The use of derivatives for managing risk can be dangerous in the absence of proper personnel, teamwork, controls, and organization. Up to now, we have largely focused on the quantitative aspects of risk management. The concerns of this chapter are mostly qualitative. It is easy to neglect these more subjective factors that lead to good risk management, but they are critically important. In fact, all of the quantitative models for and analytical knowledge about risk management would be wasted if an organization could not implement sound risk management policies.

To understand how risk management is best practiced, we must start with a perspective on the risk management industry and the profession of risk management itself. We need to understand the types of organizations that operate in the risk management industry. Then we shall take a look at how the practice of risk management fits into a typical organizational structure. We shall next turn to the very important requirement of properly accounting for risk management. Then we shall take a look at the accepted recommendations for how risk management should be practiced. Knowing that we can learn as much about how to do things well by studying how to do things badly, we shall conclude by looking at what happens when an organization does not practice good risk management.

THE STRUCTURE OF THE RISK MANAGEMENT INDUSTRY

The risk management industry is made up of end users, dealers, and other firms, including consultants and specialized software companies.

End Users

End users are generally thought of as firms that wish to engage in derivatives transactions to manage their risk. Needing to execute a derivatives transaction, end users contact derivatives dealers, whom we shall discuss in the next section. The end user could be thought of as the party who is purchasing a service while the dealer is the supplier of the service.

The end user community primarily consists of nonfinancial corporations. Other end users include investment firms, which include pension plans and mutual funds. In addition, some financial institutions that are not in the business of being derivatives dealers will still use derivatives to manage their risks. At times even a dealer, managing its own risk, will take the position of an end user opposite another dealer. Many foreign governments use derivatives. In the United States, state and local governments use derivatives, though somewhat less than corporations. With the exception of certain agencies, such as the postal service, however, the U.S. federal government does not use derivatives. Many private organizations, such as charities, endowments, and universities, use derivatives.

Obviously there are minimum firm sizes necessary to justify certain derivatives activities, but risks of as little as a few million dollars can justify over-the-counter derivatives, and even smaller sizes can be managed with exchange-traded derivatives. As one might expect, larger firms use derivatives more and have more organized and sophisticated risk management operations.

In a typical corporation the treasury department has responsibility for managing the firm's cash. It operates in the financial markets on a regular basis, borrowing cash, repaying loans, and wiring funds to and from domestic and foreign offices. If the firm engages in derivatives transactions to manage financial risk, these transactions are typically executed and managed by the treasury department. Normally a treasury department is a cost center and not a profit center, meaning that its job is to perform a function at the lowest possible cost and not to try to earn a profit. In recent years, however, many corporate treasurers, sometimes responding to pressure placed on them by senior management, have begun using derivatives to speculate and turn a profit. Unfortunately, this has put pressure on some companies to speculate in markets they know little about and has led to some excellent examples of bad risk management.

Some derivatives activities are conducted outside of the treasury. For example, a firm might have a commodity purchasing manager who hedges its exposure to commodity prices. Some firms run highly decentralized derivatives activities. They may have numerous subsidiaries, each of which engages in its own hedging. We shall discuss these types of organizational issues later in this chapter.

End users have shown a preference for particular types of derivative instruments for a given type of risk. Forward contracts are preferred for foreign exchange risk, with swaps being the primary instrument for managing interest rate risk, and futures being preferred for commodity price risk. Stock market risk is primarily managed with options, both over-the-counter and exchange-listed.

Dealers

Dealers, as previously noted, are financial institutions making markets in derivatives. These institutions include banks and numerous investment banking/brokerage firms. They quote prices indicating their willingness to take either side of derivatives transactions. These prices, as we have noted previously, are stated in terms of a bid price and an ask price, with the price they are asking being higher than the price they are bidding. Dealers will not naturally have offsetting positions. Any risk that remains is usually hedged by taking on new and offsetting derivatives, which may be in the form of exchange-traded or over-the-counter contracts. Thus, the dealer is not exposed to market movements and earns a profit off of the spread between its buying and selling prices.

Most dealer firms have highly sophisticated derivatives operations. They employ individuals with training in finance, economics, mathematics, physics, computer science, and engineering. Many of their technically trained personnel design new derivatives and risk management strategies and determine how to price them. Dealers also employ sales personnel who call on potential end users, examine their risk management needs, and try to convince the end users to engage in transactions with the dealers. Dealers generally make large investments in computer hardware and software for managing their derivatives positions. Derivatives activities have been a tremendous source of profits for many dealers.

Other Participants in the Risk Management Industry

The growth in risk management has created an industry of firms providing consulting and software services. Consulting services include general risk management consulting, legal and accounting/auditing, and personnel search. The specialized knowledge required to program computers, understand the emerging derivatives law, provide proper accounting and auditing functions, and search for specialized technical personnel has led to profitable opportunities for individuals and firms with these skills.

ORGANIZING THE RISK MANAGEMENT FUNCTION IN A COMPANY

The quality of an organization always starts at the top. This means that senior management and the board of directors must take the initiative for being knowledgeable and aware of the derivatives activities the firm engages in. This does not mean that they must be experts in derivatives, but they should be able to define each of the instruments used by the firm and should know why the firm uses them. Most importantly, they must establish written policies and procedures governing the use of derivatives. These policies must specify a rationale for the use of derivatives, define the circumstances under which derivatives can be used, authorize the appropriate personnel to execute transactions, define trading limits, establish control procedures to ensure that all policies are followed, and address the issue of how the performance of a derivatives/risk management operation will be evaluated.

Senior management is particularly responsible for ensuring that competent personnel with sufficient technical knowledge are employed by the firm when such skills are needed. They must also authorize the expenditure of funds for the necessary hardware and software.

But in order to do these things well, an organization must have a structure conducive to the proper practice of risk management. Risk management can be practiced differently depending on whether the organization is a dealer or end user.

A dealer organization, which is in the derivatives business to earn a profit, will engage in numerous transactions and naturally should practice risk management at the centralized firmwide level. If there are relationships between movements in interest rates, stock prices, and foreign exchange rates, and there nearly always are, it would be unwise to have separate risk managers monitoring the performance of the bond department, the equities department, and the various international operations. Rather, most dealers have opted for a single risk manager, who often reports to the chief executive officer. The risk manager should have access to all of the necessary financial and statistical information and should have the authority to stop traders from trading and/or force them to unwind positions. It should not be possible for traders or derivatives sales personnel to influence the risk manager. Some dealers use a risk management committee, but in either case, the risk manager's or the committee's responsibility is usually to senior management. In any case, *independent risk management is probably the most important requirement for an effective risk management system.*

A corporate or end user risk management function can be surprisingly decentralized. Some firms may have purchasing departments that hedge purchases of raw materials. Many firms will have foreign subsidiaries who hedge their subsidiary operations. Most corporations delegate domestic interest rate hedging to the treasury department. The treasury's normal responsibility is managing cash and bank loans. In recent years, many corporations have begun to centralize their risk management activities, focusing on risk management of the firm as a whole. This enables the firm to allow certain natural hedges to work. For example, a firm might be exposed long to German interest rates in one division and short in another. If each division hedges separately, unnecessary transactions are done. A centralized risk management function would determine that the risks of two divisions are at least partially offsetting.

Many corporations unfortunately do not provide for an independent risk management function. As previously noted, they sometimes operate their treasury department as a profit center. Instead of attempting to simply manage the firm's cash and bank loans efficiently, treasurers sometimes start trading in the derivatives markets with the objective of earning a profit. It is not reasonable to expect that they can do this effectively over the long run and, unfortunately, for some firms, the long run came early in the form of large derivatives losses.

Derivatives dealers have two distinct groups of derivatives specialists. One group is the sales personnel, who call on clients, analyze the clients' risks, and offer the dealer's products and services for the purpose of managing the clients' risks. The other group is made up of traders who execute derivatives transactions for the

DERIVATIVES TOOLS

Concepts, Applications, and Extensions

Professional Organizations in Risk Management: GARP and PRMIA

Financial risk management is an excellent career choice, but a difficult career to get into. As usual, experience is important, because few people can go right into a financial risk management position out of school. There are a number of schools with graduate programs in financial risk management or related aspects of financial risk management. The most important requirements, however, are a keen analytical mind, a willingness to work long hours, and a solid understanding of how financial markets work. Financial risk management jobs are found in virtually every medium-to-large bank, many large corporations, and even some nonprofit organizations and government agencies.

Practitioners of financial risk management have formed two professional organizations, the Global Association of Risk Professionals, known as GARP, and the Professional Risk Managers' International Association, known as PRMIA.

As of this writing GARP (http://www.garp.com) has about 56,000 members from around 100 countries around the world. Its objective is stated as follows:

GARP's aim is to encourage and enhance communications between risk professionals, practitioners and regulators worldwide. Through its events, publications, website and certification examination (FRM), GARP works on expanding views and increasing recognition of the global risk management community.

GARP is governed by a group of about 20 trustees and an executive committee of seven. GARP has about 24 regional chapters in about 18 countries, with each chapter having a director. The chapters conduct meetings that typically include outside speakers. GARP also holds national and international meetings and conventions. There are several levels of membership, each providing different privileges and having different levels of dues. The highest level of membership costs $100. A student membership is $50. GARP's primary focus is its professional certification program, the Financial Risk Manager (FRM) exam, a single examination offered in 40 cities around the world in mid-November. The exam consists of 140 to 150 multiple choice questions given in two 2 1/2 hour sessions. Its content is as follows:

Quantitative Analysis (10%)
Market Risk Measurement and Management (30%)
Credit Risk Measurement and Management (25%)
Operational and Integrated Risk Management (20%)
Legal, Accounting, and Ethics (15%)

The exam has been taken by about 8,400 candidates since 1997, with an overall pass rate of 51 percent. The exam is taken in booklet form, and results are provided about two months later. Early registration costs $500, with costs increasing to a maximum late registration charge of $900. GARP provides a list of study materials and sample exams.

The second organization, PRMIA (http:www. prmia.org), consists of about 35,000 members from around 169 countries as of June 2006. PRMIA's objectives are stated as follows:

The PRMIA Mission is to provide a free and open forum for the promotion of sound risk management standards and practices globally.

To accomplish our mission, we will always be:

- *a leader of industry opinion and a proponent for the risk management profession*
- *driving the integration of practice and theory and certifying the credentials of those in our profession*
- *connecting practitioners, researchers, students and others interested in the field of risk management*
- *global in our focus, promoting cross-cultural ethical standards, serving emerging as well as more developed markets*

- *transparent, nonprofit, independent, member-focused and member-driven*
- *working with other professional associations in furtherance of the PRMIA mission*

PRMIA is governed by a nine-person board of directors, an advisory panel, and various other committees. PRMIA has about 62 regional chapters. Like GARP, PRMIA conducts meetings, conferences, and conventions. PRMIA currently assesses no dues, receiving support from sponsoring organizations. PRMIA also conducts its own certification exam called the Professional Risk Manager (PRM) exam. This exam consists of four parts, and any or all of the parts can be taken on any business day. The four parts are:

Finance Theory, Financial Instruments, and Markets (30 questions, 90 minutes)
Mathematical Foundations of Risk Measurement (24 questions, 120 minutes)
Risk Management Practices (36 questions, 90 minutes)
Case Studies, PRMIA Standards of Best Practice, Conduct and Ethics (30 questions, 60 minutes)

The test is computerized, so results are known immediately. PRMIA has arranged with various testing companies that the exam can be taken in over 4,000 cities worldwide. PRMIA quotes the recent pass rate as about 51 percent. The cost is $125 for each exam, but discounts are available if more than one exam is taken on a given day. The maximum cost is $295 if the candidate takes all four exams in one sitting. PRMIA provides a recommended reading list and sample questions.

There are no clear advantages to one organization or exam over another. Anyone pursuing a career should be a member of one, if not both, organizations. Student membership is highly encouraged.

dealer. Most dealers also have research departments, normally made up of technically trained personnel who conduct research, design and price new products, and provide support for sales personnel and traders. Most corporations do not have research departments but may have a few individuals on the staff with trading experience and a high degree of technical expertise.

The derivatives operations of dealers are typically segregated into what are called the front office and back office. Some firms have an additional layer in between the front office and back office, which is naturally called the middle office. The front office consists of the traders and their assistants. The back office personnel are the clerks and operations officers whose job is to process the paperwork and generate the necessary reports. It is considered critical by nearly every dealer that these two functions be separate. If they are not, as we shall see later in this chapter in the example of the failure of Barings Bank, the proverbial "fox is guarding the chicken coop." In some firms, the back office may be the only control mechanism. While this is not the optimal structure, at the very least the back office should not have to report to the front office or feel any pressure from it.

All firms should also have a legal department or have access to external legal counsel. The attorneys are responsible for ensuring that all contracts are properly documented and conform to appropriate laws.

Within most firms there is an individual called a compliance officer who is responsible for assuring that all internal and external regulations are followed.

The accounting function is an important one in all organizations. Although an accounting system is not the same as a risk management system, the two must not be inconsistent. An accounting system also provides for periodic auditing, but does not substitute for risk management. Auditing simply determines whether the financial records have been kept in conformance with the policies established by the accounting profession and the firm. Auditing is a periodic process. Risk management is a continuous process requiring far more than checking financial entries.

Any firm that engages in derivatives and risk management activities should provide for regular evaluation of its performance. It should define its objectives and establish a system that will provide timely and unbiased assessments of the quality of its activities. This is not a simple task since objective performance evaluation is quite difficult to do, but this issue must be addressed, nonetheless.

Although there are many effective variations, Figures 16.1 and 16.2 illustrate examples of a derivatives/risk management operation for a dealer and for a corporate end user. The dealer operation is centralized with an

FIGURE 16.1 Example of a Derivatives/Risk Management Operation for a Dealer

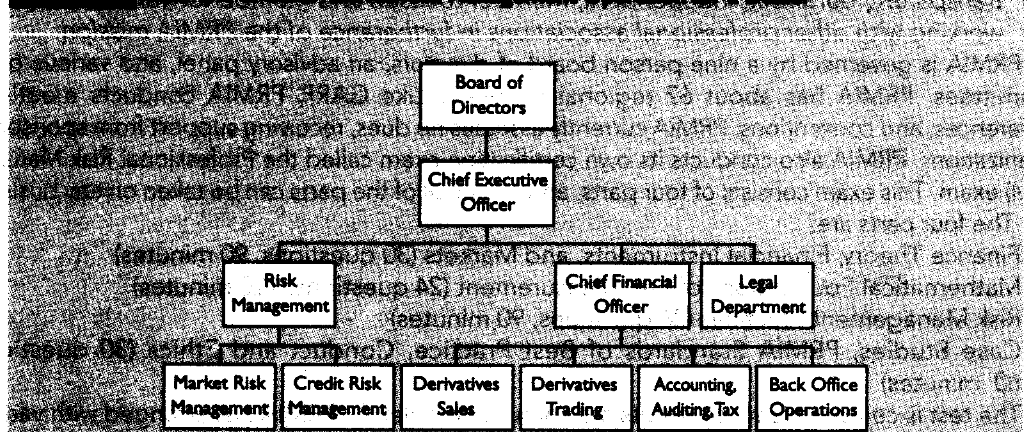

FIGURE 16.2 Example of a Derivatives/Risk Management Operation for a Corporate End User

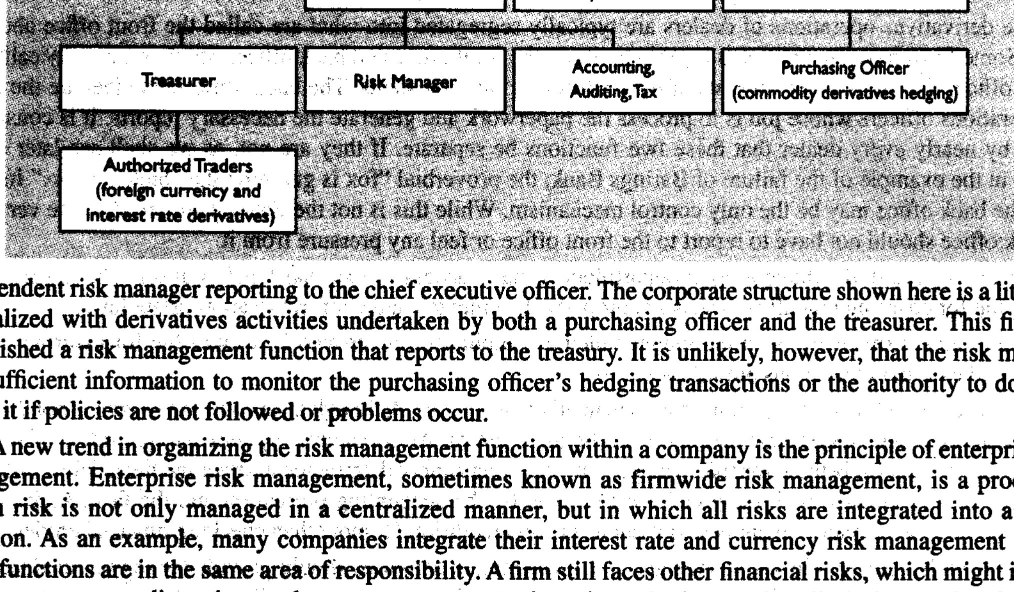

independent risk manager reporting to the chief executive officer. The corporate structure shown here is a little less centralized with derivatives activities undertaken by both a purchasing officer and the treasurer. This firm has established a risk management function that reports to the treasury. It is unlikely, however, that the risk manager has sufficient information to monitor the purchasing officer's hedging transactions or the authority to do much about it if policies are not followed or problems occur.

A new trend in organizing the risk management function within a company is the principle of enterprise risk management. Enterprise risk management, sometimes known as firmwide risk management, is a process in which risk is not only managed in a centralized manner, but in which all risks are integrated into a single function. As an example, many companies integrate their interest rate and currency risk management so that these functions are in the same area of responsibility. A firm still faces other financial risks, which might include exposure to commodity prices and even exposure to equity prices, the latter primarily in its pension fund. An integrated risk management function would bring the management of these risks under control of a single risk management function.

A true enterprise risk management system, however, would consolidate the management of other types of risks. For example, most firms spend a considerable amount of money on insurance, but the management of risks covered by traditional insurance policies is typically done in a separate function. Many risks are managed internally by efforts to avoid the exposure. For example, the risk of fire and dangerous factory work are minimized by using safety precautions. A true enterprise risk management function would bring the management of all types of risks under the same area of responsibility. In fact, the insurance industry is moving into offering financial risk management products. In some cases, financial derivatives are combined with insurance contracts to truly integrate the management of these diverse and seemingly unrelated risks.

RISK MANAGEMENT ACCOUNTING

One of the most controversial issues in the risk management world is the proper accounting treatment of derivatives transactions. Because of the unusual characteristics of many derivatives and the relative newness of many of these instruments, the accounting profession has not, until recent years, begun to catch up with the derivatives world. For many years derivatives were off-balance-sheet items, meaning that it was difficult, if not impossible, to determine from traditional financial statements what types of derivatives were being used and the effects of those derivative transactions on earnings. Much of the problem stems from the fact that the most widely used and accepted application of derivatives, hedging, leads to considerable complications in accounting. Hedges often generate cash losses and gains while the transactions they are designed to hedge generate only paper gains and losses. If derivative gains and losses are realized and reflected in income statements while gains and losses on the transactions they are designed to hedge are not realized until later, a hedge, which is designed to reduce volatility, will give the appearance through accounting statements of an increase in volatility.

One solution to the problem is to use hedge accounting. Hedge accounting permits the firm to defer the gains and losses on the derivative until later when the hedge is completed. For example, a firm that holds a stock and hedges its future sale with a derivative would not recognize the derivative gains and losses until the security is actually sold. Unfortunately, hedge accounting can lead to abuses that mislead users of financial statements. For example, suppose a firm plans to borrow money at a later date. It enters into a derivative to hedge any increase in interest rates. Suppose it incurs a gain on the derivative resulting from an increase in interest rates. Then suppose it ultimately decides not to borrow the money. It records a profit on the derivative and does not take out the loan. Had it incurred a loss on the derivative, it might well have gone ahead and borrowed, imbedding the derivatives loss in the higher interest rate on the loan. Was this even a hedge? Probably not because there was clearly no firm commitment to borrow. Yet hedge accounting was used.

Swaps, forwards, and futures pose an interesting issue in derivatives accounting. They have a zero value at the start and yet gain and lose value as the underlying moves. As the derivative gains value, it effectively becomes an asset; as it loses value, it effectively becomes a liability. Should that asset/liability be shown on the balance sheet? The practice of good accounting suggests that it should. But how?

In the United States the appropriate methods for accounting are established by the Financial Accounting Standards Board (FASB). A worldwide organization, the International Accounting Standards Board (IASB) attempts to coordinate and standardize accounting principles across countries. Neither of these bodies has any legal authority. They simply represent efforts on the part of the accounting profession to establish acceptable norms. Auditors' reports must indicate whether established accounting norms have been followed. For publicly traded companies, the securities regulators, such as the Securities and Exchange Commission in the United States, require certain standards and prohibit misleading information in accounting statements. Thus, while accounting standards technically carry no legal weight, they effectively constitute a rule of law.

In 1996 the FASB undertook a project to study and establish standards for accounting for derivatives. This culminated in the release in 1998 of Financial Accounting Standard 133, called FAS 133, *Accounting for Derivative Instruments and Hedging Activities.* This controversial document has undergone much scrutiny and criticism but has survived and appears to be the norm to which derivatives accounting will conform. Here we shall look at the important elements of FAS 133, which went into effect for most firms in June 2000. FAS 133 has been amended three times with FAS 137, 138, and 149. The International Accounting Standards Board adopted a similar document, IAS 39, which is somewhat broader in that it applies to all financial assets and liabilities while FAS 133 applies only to derivatives. IAS 39 was adopted in 2001 and implemented in January 2005.

FAS 133 basically takes the position that all derivatives transactions must be marked to market, meaning that the gains and losses incurred from derivatives transactions over the accounting period are determined, regardless of whether the derivative transaction has been terminated, and reported in financial statements. In addition, derivatives must appear on the balance sheet as assets and liabilities. But first, FAS 133 had to define what a derivative is.

Under FAS 133 a derivative has one or more underlyings and one or more notional amounts, which combine to determine the settlement or payoff. It must require either no initial outlay or an outlay that is smaller than would be required for other contracts with similar payoffs. In addition it requires a net settlement or physical delivery of the underlying.

FAS 133 permits some exceptions to the rule. For example, purchases of assets on time are not derivatives. The purchase of assets with payment to occur very shortly are not derivatives. Insurance is not a derivative. Options used in executive compensation are not derivatives, at least not with respect to the regulations prescribed by FAS 133; these instruments do, however, have their own accounting requirements. Certain imbedded derivatives may or may not be derivatives. For example, interest-only and principal-only strips are not generally derivatives, but options imbedded in certain floating-rate securities may be considered derivatives and, therefore, may have to be separated and conform to FAS 133.

FAS 133 classifies all derivatives transactions into one of four categories: fair value hedges, cash flow hedges, net investment in foreign currency operations, and speculation.

Fair Value Hedges

A fair value hedge is a transaction in which a firm hedges the market value of a position in an asset or a liability. The gain or loss on the derivative is recorded and reflected in current earnings. In addition, the asset or liability is marked to market and its gain or loss is also reflected in current earnings.

For example, a firm holds a security and enters into a derivative transaction to hedge the future sale of the security. At the end of the quarter the firm is preparing its financial statements and the hedge is still in place. Let us say that there is a loss in value of the security of $100,000.Its derivative, however, incurs a gain of $96,000. It records the following accounting entries:

Debit Derivative (for example, swap, futures, and so forth)	$96,000
Credit Unrealized Gain on Derivative	$96,000
Debit Unrealized Loss on Security	$100,000
Credit Security	$100,000

Remember that debits increase and credits decrease assets, and debits decrease and credits increase liabilities and capital. Credits increase and debits decrease income. Thus, in the first entry, an asset account (the derivative) is increased and income (the unrealized gain) is increased. In the second entry, income is decreased (the unrealized loss) and an asset (the security) is decreased. The net effect is that both income and assets are decreased by $4,000.

If the hedge is effective, it will have little effect on income and, therefore, on earnings. Otherwise, earnings will fluctuate. Hence, effective and ineffective hedging will be properly reflected in earnings.

To qualify for fair value accounting, firms must maintain proper documentation of every transaction and the hedge must be expected to be effective. The sale of options, however, is ordinarily not eligible for fair value accounting. In addition, bonds that are expected to be held to maturity are not eligible for hedge accounting. Certain anticipated transactions, however, can be eligible for hedge accounting, though many of these will fall under the following category.

Cash Flow Hedges

In a cash flow hedge, the risk of a future cash flow is hedged with a derivative. A classic example is a firm planning to take out a loan. It might sell Eurodollar futures or buy an FRA to hedge the risk of an increase in interest rates. Cash flow hedging is somewhat more complicated than fair value hedging and must meet a more stringent requirement.