# CHAPTER 20

# Current Topics in Risk Management

In this chapter, we discuss several currently important risk management topics: value at risk (VaR), credit derivatives, options on debt instruments, swaptions, and exotic options.

## 20.1 VALUE AT RISK (VAR)

### 20.1.1 Background

One goal of active risk management is to reduce the variability of uncertain cash flows. As you have seen, however, risk management cannot eliminate this variability. Further, many examples that we have presented in this textbook have presented risk management vehicles dealing with a single risk, such as interest rate risk, foreign currency risk, commodity price risk, or stock market risk. However, the modern corporation could have hundreds, or thousands, of sources of uncertainty that are being hedged (or not being hedged).

Value at risk (VaR) is the name of a risk management concept by means of which senior management can be informed, via a single number, of the short-term price risk faced by the firm. The origin of "value at risk" stems from a request by J. P. Morgan's chairman, Dennis Weatherstone, for a simple report, to be made available to him every day, concerning the firm's risk exposure. Since then, VaR has rapidly become the financial industry's standard for measuring exposure to financial price risks. Today, few financial firms fail to make VaR part of their daily reporting to senior management.

Use of the VaR concept has become pervasive.[1] The Bank for International Settlements (BIS), which is essentially the central bank of the world's major central banks, proposed in April 1995 that major banks use VaR to determine their capital adequacy requirements. These requirements became effective January 1, 1998. The U.S. Federal Reserve Bank and the International Swaps and Derivatives Association (ISDA) basically endorsed the BIS's recommendations about VaR. In December 1995, the U.S. Securities and Exchange Commission (SEC) proposed rules that would require corporations to disclose information concerning their use of derivatives. Firms would be directed to use one of three methods that would provide information about the risk exposure of their portfolios of financial assets and derivatives. One of the methods was VaR.[2] In April 1996, eleven individuals from the institutional investment community formed the Risk Standards Working Group (1996) and charged themselves to "create a set of risk standards for institutional investment managers and institutional investors." Risk Standard 12 states that money managers "should regularly measure relevant risks and quantify the key drivers of risk and return." The standard proceeds to suggest VaR as one possible method for measurement of risk.[3]

The website www.gloriamundi.org/ is devoted to the concept of VaR.

### 20.1.2 The VaR Concept

The VaR concept is an attempt to encapsulate an estimate of the price risk possessed by a portfolio of derivatives and other financial assets. The statistic that VaR provides is the dollar amount by which the value of a portfolio might change with a stated probability during a stated time horizon; the time horizon might be one day, one week, or longer. For example, a financial institution might estimate that there is a 1% probability that its portfolio will decline by more $15 million during the next week. The decline in value might be caused by changes in the prices of fundamental risk factors such as changes in foreign exchange rates, interest rate changes, commodity price fluctuations, changes in stock prices, and/or increases in volatility.

It is important to note that the price risk number obtained from a VaR model summarizes risk exposure into a dollar figure that purportedly represents the estimated maximum loss over an interval of time. That is, a dollar loss greater than the VaR estimate will occur with a smaller probability than the VaR estimate.

### 20.1.3 Computing VaR

There is no standard method to compute VaR. However, several accepted methods for computing VaR have emerged. These methods are the variance–covariance approach, the historical simulation method, and the Monte Carlo simulation method. It is important to note that VaR is sensitive to the assumptions made and to the method used. In other words, your VaR number will differ depending on your assumptions and which method you have chosen to compute it. Risk managers should understand both these sources of sensitivity when choosing a method to compute VaR.

#### 20.1.3.1 The Variance–Covariance Approach

The **variance–covariance approach** is also known as the "delta-normal method." The variance–covariance approach is used by J. P. Morgan's *RiskMetrics* model.[4]

The important assumption of this approach to estimate VaR is that the returns for each of the institution's assets are normally distributed. Recall from your investments course that the variance of a portfolio's returns are computed using the following formula:

$$\text{Var}(\tilde{R}) = \sum_i \sum_j w_i w_j \sigma_{ij} = \sum_i \sum_j w_i w_j \sigma_i \sigma_j \rho_{ij} \tag{20.1}$$

where

$w_i$ = fraction of the total portfolio value consisting of asset $i$, where $\Sigma w_i = 1.0$
$\sigma_{ij}$ = covariance of asset $i$'s returns with asset $j$'s returns
$\sigma_i$ = standard deviation of asset $i$'s returns
$\rho_{ij}$ = correlation of asset $i$'s returns with asset $j$'s returns

For two assets, equation (20.1) becomes

$$\text{Var}(\tilde{R}) = w_1^2\sigma_1^2 + w_2^2\sigma_2^2 + 2w_1w_2\sigma_{12} = w_1^2\sigma_1^2 + w_2^2\sigma_2^2 + 2w_1w_2\sigma_1\sigma_2\rho_{12} \tag{20.2}$$

Suppose the risk manager assumes that the return distribution of his portfolio is normal with a mean of zero. Then, the first step to obtain a VaR estimate is to obtain an estimate of the variance of his portfolio's periodic returns. These periodic returns could be calculated daily, weekly, or for any other interval.

After one has obtained a portfolio mean return and return variance, it becomes a simple statistical procedure to estimate what the loss in value of the portfolio will be during that period. Further note that this estimate can be calculated with any desired probability. For example, the value at risk, VaR, will equal 1.645 times the portfolio's standard deviation with a probability of 5%. The maximum loss with a 1% probability will equal 2.327 times the portfolio's standard deviation.

In one approach for using the variance–covariance method to estimate VaR, the variance of each asset's returns and each pairwise covariance are estimated. These variance estimates can be calculated from historical data, the implied volatilities contained in option prices, and/or volatility forecasts, or they can be generated from the risk manager's subjective beliefs. If the risk manager has only two assets in his portfolio, $w_1$ and $w_2$ in Equation (20.2) are the fractions invested in each asset, $\sigma_1$ and $\sigma_2$ are the standard deviations of the returns of each asset, and $\rho_{12}$ is the correlation between the returns of the two assets.

In practice, many VaR users employ changes in value instead of using rates of return. That is, let $V_t$ represent the value of an asset on day $t$. Then, the change in value is simply $V_2 - V_1$. The percentage change in value (i.e., the rate of return) is given by $(V_2 - V_1)/V_1$.

There are a few reasons to use value changes instead of percentage changes in value. First, many derivatives, such as futures contracts, are marked to market. In these cases, their values are zero after they have been marked to market, and the percentage rate of return cannot be computed. Second, many derivatives will have a negative value; that is, they are liabilities. Regardless of whether a contract is an asset or a liability, the user of VaR is interested in the change of the value of the contract over the period of interest. Thus, the standard deviations of value changes are of greater use than the standard deviations of rates of returns.

---

**EXAMPLE 20.1** Using the Variance–Covariance Approach to Calculate VaR: Consider a portfolio consists of these three assets:

**a.** *A currency swap.* Because of changes in the exchange rate since the swap was first entered into, the swap now has a value of $2 million, or 8.7% of the portfolio's total value.

**b.** *A bond.* The market value of the bond is $17 million, which is 73.9% of the portfolio's total value.

**c.** *A stock.* The 10,000 shares are worth $4 million, or 17.4% of the portfolio's total value.

Assume the variance–covariance matrix[5] of the assets' daily returns is

| | Swap | Bond | Stock |
|---|---|---|---|
| Swap | 0.00090 | –0.00008 | 0.00007 |
| Bond | | 0.00040 | –0.00010 |
| Stock | | | 0.00300 |

Then the variance of the portfolio's daily returns distribution is then found by using Equation (20.1):

$$\begin{aligned}\text{Var}(\tilde{R}) &= w_1^2\sigma_1^2 + w_2^2\sigma_2^2 + w_3^2\sigma_3^2 + 2w_1w_2\sigma_{12} + 2w_1w_3\sigma_{13} + 2w_2w_3\sigma_{23} \\ &= (0.087)^2(0.00090) + (0.739)^2(0.00040) + (0.174)^2(0.00300) \\ &\quad + 2\,(0.087)(0.739)(-0.00008) + 2(0.087)\,(0.174)(0.00007) \\ &\quad + 2\,(0.174)(0.739)(-0.00010) \\ &= 0.0002822\end{aligned}$$

The standard deviation of the daily returns distribution is $(0.0002822)^{1/2} = 0.0168$, or 1.68%. One standard deviation of dollar loss from the portfolio value of \$23 million is \$386,375 (1.68% of \$23 million is \$386,375). To calculate the VaR, multiply the number of standard deviations for a stated probability level by one standard deviation of dollar loss. Thus, there is a 5% probability that a one-day loss of (1.645)(\$386,375) = \$635,587 will be realized. There is a 1% probability that a one-day loss of (2.327)(\$386,375) = \$899,095 will be realized.

An important assumption we made was that returns, or value changes, are normally distributed. Option returns are, of course, not normally distributed. The lack of normality becomes most obvious when options are held until expiration, as the returns distribution of a long put or long call is truncated at −100% and there is a great deal of positive skewness.

One way to avoid the "nonlinear" payoff pattern of options is to assume that the options are equivalent to delta units of the underlying asset. Recall from Chapter 16, that the BOPM teaches us that a call is equivalent to owning delta shares of stock and borrowing. Thus, the risk manager can substitute $\Delta S + B$ for any option. The returns distributions for stocks and bonds are closer to normal than are those for options. Remember, however, that the concept of delta is applicable only when all else is held equal. Thus, substituting $\Delta S + B$ for an option works best when one is calculating a daily VaR. For periods longer than one day, the time to expiration is definitely not being held constant, and it is unlikely that the riskless interest rate and the implied volatility of the underlying asset will remain constant. Option deltas change when any fundamental determinant of option value changes, such as time to expiration. Moreover, deltas are unstable when gammas are high.

Another approach for using the variance–covariance approach to estimate VaR requires the decomposition of asset value volatility into its most fundamental sources. This is the approach taken by J.P. Morgan's *RiskMetrics* model. The fundamental factors behind the changes in asset values are interest rates in each relevant country and for several times to maturity, spot exchange rates, country stock market factors, and commodity prices. Think of a forward exchange contract. The basic factors that cause it to change in value are the domestic interest rate for a zero-coupon bond maturing on the same day as the forward contract's delivery date, the foreign interest rate for the same time to maturity, and the spot exchange rate.

There are several potential problems with the variance–covariance approach for estimating VaR. First, if historical data are used to estimate the variance–covariance matrix of asset returns, it is highly likely that the future will not replicate the past. In particular, the estimation period might not include an "event" that happens only once every '*N*' years, such as the 1987 stock market crash or the 1997 collapse of the Asian stock markets and currency values. Similarly, the recent past may actually include a period in which an "event" occurred that is unlikely to ever occur again. Then, the probability of a large loss will be overstated.

VaR figures that are calculated from historical data will differ, depending on how far back in the past the historical data are gathered. If one year of daily historical data were gathered on January 1, 1989, these data would not include the very unusual financial events of October 1987. But if two years of data are gathered, the estimated variances and covariances are very different from those estimated using only one year of data from 1988 only. Consider a $23 million portfolio that mimics the S&P 500 Index. The mean daily return and standard deviation of daily returns for the S&P 500 Index were as follows:

| | **1988 Only** | **1987 and 1988** |
|---|---|---|
| Mean | 0.000284 | 0.000237 |
| Standard deviation | 0.010613 | 0.016858 |
| 5% VaR | $401,543 | $637,822 |
| 1% VaR | $568,018 | $902,257 |

As you can see, the effect on the VaR calculation is dramatic. For example, assuming a mean of zero, if one were using 1988 data only to calculate VaR on January 1, 1989, the 1% VaR is $568,018. However, if two years of daily data are used, the 1% VaR is $902,257, about 60% higher. Recall that the purpose of VaR is to encapsulate risk into one number so that senior management can ascertain the firm's price risk. Senior risk managers must be aware of this type of data effect on VaR.

Another shortcoming that is somewhat related to the foregoing discussion is that return distributions may not be normal. The return distribution might be skewed, or it may possess what is called "fat tails." A fat-tailed distribution is said to possess leptokurtosis and is characterized by too many observations (relative to a normal distribution) in the tails. That is, unusual events occur more frequently than a normal distribution would predict.

Finally, as we have already noted, options, and assets with optionlike characteristics, possess nonnormal return distributions. Their sensitivities to changes in their determinants of value ($S$, $r$, $\sigma$, and $T$) are themselves unstable and unpredictable. Moreover, their sensitivities (e.g., their deltas) are estimated only for infinitesimal changes in their fundamental determinants of value.

### 20.1.3.2 The Historical Simulation Approach

In the historical simulation approach, the risk manager determines how each relevant price has changed during each of the past '*N*' time periods. For example, the relevant prices might be the dollar price of the Japanese yen and euro, short-term U.S. interest rates, and the dollar price of oil. It is possible to use a historical database to find the price changes for each of these four variables on each of the last 300 days. The risk manager then estimates how each of the 300 sets of price changes would affect the value of his current portfolio of spot assets and derivatives. Assuming that each of these outcomes is equally likely, he will then rank the resulting estimated value changes from most positive (value increases) to most negative (value decreases). If the risk manager is interested in the VaR at the 95% confidence level, he will find the 15th worst outcome out of the 300. The VaR at the 99% confidence level is the 3rd worst outcome. These are the losses that will be exceeded only 5% of the time and 1% of the time, respectively.

Institutions might be interested in their VaR for the next day, or for longer intervals of time. The more actively the portfolio is traded, the shorter the time interval of interest. If the portfolio is turned over frequently, then estimating the monthly VaR for today's portfolio provides little information. Thus, tomorrow's portfolio is likely to be very different from today's portfolio.

Institutions might also wisn to use the last 100, 250, and 500 days to estimate their daily VaR. Note that there is no assurance that any given amount of historical data is "correct." The risk manager must decide whether the distant past or the recent past more accurately predicts what tomorrow will bring.

It is important to note that using the historical simulation approach implicitly incorporates the correlations among asset price changes. To the extent that, say, price changes of the yen and the euro are highly correlated, this correlation will show up in the historical set of price changes.

#### 20.1.3.3 The Monte Carlo Simulation Approach

In the Monte Carlo simulation approach, the risk manager specifies probability distributions or stochastic processes for prices in the future. A different probability distribution can be assumed for each pricing determinant. For example, changes in interest rates might be skewed right, while changes in oil prices might be drawn from a uniform distribution and changes in the price of a currency might be normally distributed.

After the distributions or processes are defined, random realizations of outcomes can be simulated. Each randomly chosen outcome is a set of prices. Correlations are explicitly incorporated in the simulation. The change in value of each asset in the portfolio is estimated for each randomly selected set of prices, thereby producing a probability distribution of future value changes of the portfolio. The risk manager can then determine what the worst outcome will be with a desired confidence level.

### 20.1.4 Stress Testing

For any VaR calculation method, it is important for the risk manager to perform *stress testing*. The purpose of stress testing is to assess the valuation impact of worst-case scenarios, regardless of whether these outcomes were realized during the recent past.

One benefit of stress testing is that economic variables can be ranked in order of their effect on portfolio value. If, say, the risk manager discovers that his portfolio's value is most sensitive to a rise in short-term U.S. interest rates, he could decide to hedge that price risk specifically. Stress testing also allows the risk manager to examine the effect of economic conditions other than price changes. For example, changes in the shape of one country's yield curve, changes in options' implied volatilities, or changes in correlations might be important VaR determinants.

#### 20.1.4.1 Stress Testing: The Importance of Skewness and Leptokurtosis

It is well recognized that financial asset returns are not distributed normally. are widely known In fact, the distributions of asset returns exhibit skewness and leptokurtosis. These exotic-sounding terms are merely ways of describing how asset returns are distributed.

Measures of central tendency and dispersion (i.e., the mean and the variance) are widely known. In statistical terms, the expected value, or mean, is known as the first moment of a probability

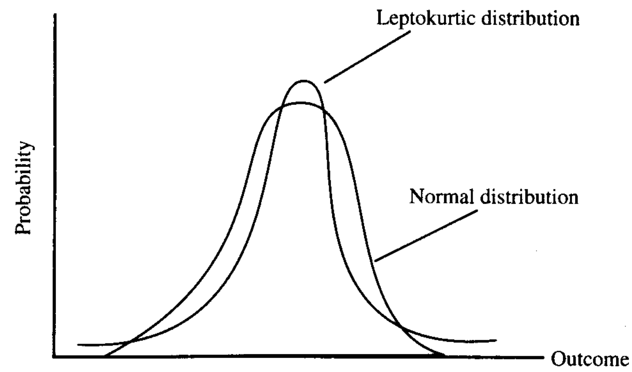


**Figure 20.1** Normal distribution and leptokurtic distribution. Relative to the normal distribution, the leptokurtic distribution has a greater probability that an observation near the mean will occur (the distribution in more peaked) and a greater probability that observations far from the mean will occur (the distribution possesses fat tails).

distribution and the variance is the second moment. Skewness, a measure of symmetry, is known as the third moment, and kurtosis is known as the fourth moment. Kurtosis can best be described in terms of the well-known symmetric, normal distribution. In comparison to a normal distribution, a distribution that possesses *leptokurtosis* is one that has too many observations near the mean and too many observations far from the mean (in the so-called tails of the distribution). Figure 20.1 illustrates the difference between a normal distribution and a distribution that exhibits leptokurtosis.

The risk manager must realize that financial asset returns display leptokurtosis and build this fact into a stress testing system. If returns are normally distributed, the probability of the event occurring might be 0.000001, or one in a million. But, because actual returns distributions are leptokurtic, or fat-tailed, low probability events occur more frequently than predicted by normal distributions.

### 20.1.4.2 The 1987 Stock Market Crash: A Lesson in Leptokurtosis

Imagine the following scenario. It is October 1, 1987. Suppose an extremely creative risk-managing group calculates the mean and standard deviation of the daily returns on the S&P 500 Index to calculate a portfolio VaR. They find the following.

| Time Period | Mean | Standard Deviation |
|---|---|---|
| 1/1/1987–9/30/1987 | 0.001505 | 0.009840 |
| 1/1/1986–9/30/1987 | 0.000952 | 0.009538 |
| 1/1/1983–9/30/1987 | 0.000690 | 0.008385 |
| 1/1/1950–9/30/1987 | 0.000309 | 0.007841 |

In addition, for the period between January 1, 1950, and September 30, 1987, the risk-managing team has uncovered the following information. Of the 9571 trading days, the market was down on 4430 days. Because the members of the risk managing team are concerned about down days, they rank the returns from lowest to highest, which allows them to prepare the following table.

| Market Falls by | Number of Days | Percentage of Down Days |
|---|---|---|
| >7% | 0 | 0.0000 |
| 6–7% | 2 | 0.0451 |
| 5–6% | 3 | 0.0677 |
| 4–5% | 6 | 0.1354 |
| 3–4% | 19 | 0.4289 |
| 2–3% | 103 | 2.3251 |
| 1–2% | 757 | 17.0880 |

The risk management team also finds that there are 222 days when the market fell by more than 1.60% (4430 × 0.05) and 44 days when the market fell by more than 2.50% (4430 × 0.01). The three worst one-day losses occurred on May 28, 1962, September 26, 1955, and June 26, 1950. However, the fourth worst one-day loss (–4.9%) occurred on September 11, 1986.

Armed with the data, the risk-managing team uses the time period January 1, 1950, through September 30, 1987 to calculate the one-day VaR for the firm's $100 million portfolio.

| VaR Calculation Method | VaR Level | VaR in Dollars |
|---|---|---|
| Variance–covariance | 5% | –$1,289,845 |
| Variance–covariance | 1% | –$1,824,600 |
| Historical simulation | 5% | –$1,600,000 |
| Historical simulation | 1% | –$2,500,000 |

After filing their report with senior management, the risk management team embarks on its well-deserved one-month vacation. Lolling on the beach on October 19, 1987, they hear that the market fell that day by 22.90%. The team is stunned. Their diversified portfolio has suffered a terrific blow. They quickly convene and ask themselves the following question:

> If financial asset returns are normally distributed, what is the probability of observing a market crash like the one that occurred on October 19, 1987?

Because the daily mean return is not statistically different from zero, the 22.90% decline represents

$$\frac{22.90}{0.7841} = 29.205$$

standard deviations from the mean. If asset returns are normally distributed, statistically speaking, the 1987 stock market crash essentially should not have occurred. In fact, given the foregoing data, if asset returns are normally distributed, 4% market declines should occur about once in every 5.9 million trading days. Yet, there were 11 days in 9571 with a market decline of 4% or more. Events like the 1987 stock market crash defy probability, yet they do occur. Further, it is quite possible that a similar event will occur again in your lifetime, because stock return distributions possess leptokurtosis.

It is important to recognize that the value-at-risk concept is designed for "normal" market conditions. However, sometimes the risk management team is interested in analyzing the performance of the portfolio under extreme market conditions. This demand has given rise to *CrashMetrics*, a data set and methodology that can be used to estimate the exposure of a portfolio to extreme market movements. Basically, the *CrashMetrics* method assumes that the effects of a crash cannot be hedged and then proceeds to predict the worst outcome for the portfolio.[6]

### 20.1.5 Back Testing

Because of the statistically perverse nature of asset returns, it is important that risk managers perform *back testing*. In back testing, the risk management team examines the performance of their VaR estimates of extreme losses with respect to realized losses. That is, back testing allows the risk manager to determine whether the VaR methods employed are adequate.

When back testing, the risk manager must be aware that there will be periods in which actual losses will exceed those predicted by VaR. For example, the risk manager must realize that statistically speaking, actual losses will exceed a 5% VaR, 5% of the time. Over a 250-day trading year, this will occur on about 12 days (i.e., $0.05 \times 250 = 12.50$).

On the days that actual losses exceed VaR predictions, the risk management team must perform a postmortem of sorts. What went wrong? Was the VaR employed appropriate? Was it defective? Or, was it "one of those days" when economic events simply overpower the portfolio? Regardless, the risk manager must remain vigilant and ask whether the extreme events merely represent an extreme, or whether they signal that the return distribution assumption must be modified. Of particular importance to the risk manager is the degree to which the VaR is exceeded. That is, it is one thing to post actual losses just a shade lower than a 5% VaR but something quite different to find that actual losses are 1.5 times the prediction by a 5% VaR.

Sometimes, the composition of the portfolio can drive actual losses beyond VaR. If selling an asset in one day can be accomplished only by accepting a large price discount, the value change caused by an adverse set of price changes should reflect this. Accordingly, bid prices should be used for the computation of VaR, particularly if the risk manager believes that parts of the portfolio will be liquidated after adverse price movements. For this reason, institutions will often adjust their VaR for the liquidity of their positions.

### 20.1.6 Which VaR Calculation Is Best?

Unfortunately, this obvious question has no simple answer. Each institution has a portfolio with some unique characteristics. Each VaR calculation method has its strengths and weaknesses. Under some conditions the variance–covariance model is superior, in other instances the historical simulation is quite handy; and the Monte Carlo simulation technique is the best for complex cases. Thus the risk manager must be familiar with all three methods.[7]

## 20.2 Credit Derivatives and Options on Debt Instruments

### 20.2.1 Credit Derivatives Background

Credit risk permeates market economies. Credit risk incorporates several types of "credit events," including changes in credit spreads or relative prices, changes in bond ratings or credit ratings, and outright default. As the chance of default increases, the prices of an issuer's bonds decline, and their yields to maturity rise. It is also true that the interest rate that must be paid by bond issuers on any floating rate debt and newly issued debt will also rise. However, in no way do increased interest payments protect the bondholder from the chance of default by the bond issuer. Indeed, higher coupon payments actually increase the probability of default.

Because changes in credit worthiness, most notably default, can result in material losses, financial institutions look to protect themselves against credit risk. Traditionally, these firms have tried to construct a portfolio composed of geographically diversified low risk loans. However, a complete shifting of default risk generally cannot be accomplished through portfolio selection. Credit derivatives, which are also contracts, have emerged to fulfill the important economic function of shifting credit risk. Credit derivatives first appeared in 1991 and have grown into an important component of the risk management function. By the end of 2001, the size of the credit derivatives market was estimated to be as much as $1 trillion.[8]

The risk of default is difficult to measure. However, in the spirit of VaR, J. P. Morgan created *CreditMetrics*, which is composed of data sets and models that attempt to help measure default risk.[9]

### 20.2.2 Major Types of Credit Derivative

There are many different types of credit derivative and more are being invented each year. Thus, it is impossible to describe every possible credit derivative. A credit derivative has payments made contingent on a well-defined credit event: changes in credit spreads or relative prices (hence periodic rates of return), changes in bond ratings or credit ratings, or default. Credit derivatives exist as forward contracts, swaps and options; all are OTC contracts. The underlying debit instruments (also called the "reference credit") might be a sovereign (i.e., made by a foreign government) loan or bond, a corporate loan or bond, or a bank loan. Often, the reference credit is a portfolio of these credits. Credit derivatives might be cash settled, or settled with physical delivery of the bond or loan.

A risk manger should be aware of the two types of credit derivatives described below.

#### 20.2.2.1 Total Return Swaps

A **total return swap** is a widely used credit derivative. In a common form of a total return swap, the party wishing to shift default risk agrees to receive a floating riskless reference rate, perhaps based on LIBOR or a Treasury security, and to pay the total return on a debt instrument.[10] The buyer of credit risk receives the return on the risky underlying credit and pays the rate based on LIBOR.

The total return on the risky debt instrument can be decomposed into a component identifiable as a return for bearing interest rate risk and a return for bearing credit risk. Because the swap can be structured to eliminate interest rate risk through the variable riskless interest payments, a total return swap can isolate credit risk.

As an example, suppose a manager of a bond portfolio want to reduce his exposure to a decline in the credit worthiness of MissMolly.com. The portfolio manager owns a bond issued by MissMolly.com, which is selling at par, matures in 10 years, and has a coupon payment.[11] Instead of simply selling the bonds in his portfolio, the fund manager decides to enter into a total return swap with a 24-month tenor. Most credit derivatives have a tenor that is less than the time to maturity of the underlying asset.

Under the terms of the swap agreement, the fund manager agrees to receive a floating payment of LIBOR + 125 basis points and pay the swap dealer the total return on the bond (i.e., the coupon and capital gain return). Payments are made every six months, and the notional principal is set equal to the fund's investment in Miss Molly.com bonds. Note that if the total return on the bonds is negative due to a rating downgrade, the total return payer (i.e., the bond fund manager) will actually *receive* two payments: the negative return on the risky bond plus LIBOR + 125 bp. If the total return on the risky bond is positive, the bond portfolio manager pays that return to the swap dealer and receives LIBOR + 125 bp.

Thus, you can see that the bond fund manager is protected from a decline in the credit quality of the issuer. Without the swap in place, the bond fund manager would have suffered major losses due to the loss in bond value caused by the rating downgrade.

In another feature of some total return swaps, if the bond issuer defaults, the total return payer can deliver the bond and receive its face value in return.

#### 20.2.2.2 Credit Default Swaps

Suppose a bank has lent money to a major corporation and fears that the borrower will default. This bank could purchase a credit default swap. A credit swap is really an option. Under the terms of a credit swap, the credit swap buyer (i.e., the bank) pays a lump sum[12] to the credit swap seller. The credit swap buyer has shifted the risk of default of the bond in question to the credit swap seller. If there is a default later during the credit swap tenor, the credit swap buyer will receive a default payment for the credit swap seller. If there is no default, the bank will receive nothing. Unlike traditional swaps, the credit buyer receives a payment only if there is a default.

### 20.2.3 Other "Event" Derivatives

Traditionally, derivative securities were designed to shift price risk. Further, this price risk was traditionally thought to evolve gradually. That is, even though dramatic price moves were sometimes encountered, the security was not designed around a catastrophic event. However, as already discussed, derivative securities have evolved to shift risk occurring from a dramatic set of circumstances, such as a "credit event," like a default.

The OTC credit derivative market is not the only market in which the risk of dramatic events can be hedged. Two innovative contracts were developed at the CBOT in corn yield insurance futures and catastrophe insurance options, although both failed. At the CME, weather derivatives now trade.

The purpose of the corn yield insurance futures was to allow farmers to hedge against quantity risk. Coupled with the CBOT corn futures contracts that allow price risk reduction, these new contracts attempted to permit further overall risk reduction for corn growing enterprises in terms of both price and quantity (the elements of total revenue).

Catastrophe insurance options were designed to allow insurance companies to shift the dollar losses resulting from such natural disasters such as tornadoes, floods, earthquakes, and hurricanes. The underlying assets for these options were a set of "loss indexes" published daily by Property Claim Services (PCS). Each loss index tracks PCS estimates for insured industry losses resulting from catastrophic events.[13]

A significant portion of the U.S. economy is directly affected by weather. Because the Chicago Mercantile Exchange lists Weather Futures and Options, firms are no longer obliged to stand by and let the weather dictate revenues, costs, and profits. These contracts are based on indexes of heating degree-days and cooling degree-days. A degree-day is defined as how much a day's temperature deviates from the benchmark of 65 degrees. Thus, these derivative contracts are designed to shift the risk of adverse temperatures (either too hot or too cold). The Merc's website, (www.cme.com/products/index/weather/products_index_weather.cfm) contains a full description of these new exchange-traded derivatives. Many private companies also offer weather derivatives. These include Koch Industries (www.entergykoch.com/index.asp) and Aquila (www.aquila.com/solutions/weather/).

Finally, degree-days have a direct impact on electricity demand and, as such, on electricity prices. To hedge against fluctuations in electricity prices, the NYMEX and IPE introduced electricity futures and options. A remarkable aspect of electricity is that it is not storable. Thus, the cost-of-carry relationship (which is based on storage) does not have to hold for electricity prices. This means that electricity traders have to face a whole new set of pricing factors. For more information on these markets, see the NYMEX (www.nymex.com) and IPE (www.IPEmarkets.com) website.

### 20.2.4 Options on Debt Instruments and Interest Rate Options

Options on debt instruments and interest rate options differ significantly from options on equities. Because they can be thought of in any of four ways, options on debt instruments and interest rate options can be ambiguous contracts.

**1.** *They are options on cash debt instruments such as spot T-bills, T-notes, or T-bonds.* However, under this interpretation, one must confront the fact that as time passes, there will be changes in the nature of the underlying asset; that is, the security's time to maturity will change. As this occurs, the interest rate that is appropriate for discounting the debt instrument's cash flows will change. For example, consider a call option today on a pure discount bond that has six years to maturity. Its price equals the present value of its face value, where the discount rate is the six-year rate. One year hence, the appropriate discount rate will be the five-year rate. Today's six-year rate will likely differ from next year's five-year rate. In addition, we can see from the appendix to Chapter 9 that longer duration bonds have greater price volatility than shorter duration bonds. Therefore, as the underlying asset's time to maturity declines, its price volatility will also decline.

**2.** *They are options on forward contracts on the underlying asset.* This characterization is definitely appropriate for European options on debt instruments, or for situations in which early exercise of American options is unlikely. The underlying asset is what will exist on the delivery date. That is, if a European call option on a spot six-year bond has one year to expiration, it is essentially a call option on a forward five-year bond.

**3.** *They are options on spot interest rates.* The spot interest rate is the one that discounts the cash flows of the cash debt instrument described in item 1.

**4.** *They are options on forward interest rates.* The appropriate rate is then the forward discount rate appropriate for the forward security in item 2.

Standardized options on spot T-bills and specific T-notes were once traded on the AMEX, However, because these contracts did not generate sufficient volume, they failed. Options on specific T-bonds, and on interest rates themselves, continue to trade on the CBOE, in what must be described as very thin markets.

One reason for the failure of options on long-term spot debt instruments as viable trading contracts, while options on long-term interest rate futures have succeeded, is that there is a limited supply of any given T-bond or T-note. Because of this limited supply, traders who control the available supply of a particular T-bond or T-note could possibly manipulate the option price on a specific security. In contrast, the underlying asset behind CBOT Treasury note and Treasury bond futures options are the futures contracts themselves.

Potential problems will arise when the liquidity of one particular debt security is poor. An option contract requires a liquid underlying asset. Since most futures contracts are more liquid than any specific spot debt instrument, options on Treasury futures have thrived, while options on the spot debt instruments themselves have not.[14]

A tremendous demand has emerged for custom-made interest rate options, such as caps, floors, and collars. Banks and other financial institutions offer these products to their clients, who either have portfolios of debt securities or have borrowed money (frequently at variable rates). The banks' customers often desire the payoff patterns offered by options or combinations of options such as spreads. Government securities dealers also buy and sell specialized debt instrument option contracts. The banks that buy and/or sell these options protect themselves by using exchange-traded options, or by the dynamic option replication techniques covered in Chapter 17.

In the next section, we describe some of the varied approaches that might be used to price options on debt instruments and interest rates.

#### 20.2.4.1 Why the Black–Scholes Option Pricing Model Is Not Really Appropriate for Pricing Options on Debt Instruments and Interest Rates

As a first approximation, the Black-Scholes option pricing model can be used to estimate the value of European calls and puts on debt instruments. Let $S$ represent the price of the cash bond. Therefore, $S$ is the present value of the future cash flows to the bondholder, discounted at the appropriate discount rate, $R$. The cash flows include the coupon payments and principal.

$$S = \sum_{t=1}^{T} \frac{CF_t}{(1+R)^t} \tag{20.3}$$

Then, $S$ is the price of the underlying asset in the BSOPM:

$$C = SN(d_1) - Ke^{-rT} N(d_2)$$

where all the parameters are as defined in Chapter 17.

Because of the severe violations of the BSOPM assumptions, however, the BSOPM is not well specified for the bond as an underlying asset. Specifically, these violations include the following.

**1.** Most bonds pay semiannual coupons. The true value of the underlying asset is the total price of the bond, which equals the quoted price plus accrued interest. This value will be discontinuous around the coupon payment date, which means that the bond's spot price cannot follow a continuous diffusion process. Still, we can make adjustments similar to those that were made to handle discrete dividend payments in Chapter 17. That is, we cannot reduce the total spot price (today's quoted price plus today's accrued interest) by the present value of coupons and accrued interest between today and the option's expiration date. This is appropriate for European options, and effectively makes the underlying asset a forward bond. If a forward price is readily available, just use that price in the BSOPM as the underlying asset. This also means that Black's model [Equations (18.15) for calls and (18.16) for puts] should be used to value interest rate options.

Summarizing, use $S^* = S + AI - PV(C)$ in the BSOPM as the underlying asset. The total price of the bond (quoted price plus accrued interest) is $S+AI$, and $PV(C)$ is the present value of any coupons between today and the option's expiration date.

**2.** The price of a cash debt instrument will not follow the diffusion process assumed by the BSOPM because a bond's price is always drawn to its face value. As time passes, there will always be a force that brings discount bond prices up to face value, even if interest rates rise. The same force draws the prices of premium bonds down to face value, even if interest rates decline. At maturity, a bond's price equals its face value.

**3.** If bond *prices* followed a diffusion process, negative yields to maturity might be implied at sufficiently high prices. Negative bond yields are highly unlikely to exist, so there must be bounds on bond prices that preclude such yields from occurring. But with the imposition of such limits, the bond price will no longer follow a diffusion process that allows the use of the BSOPM to value an option on the bond. If you assume that interest rates follow a diffusion process, the interest rate becomes the underlying asset of the model. Better still, it could be argued that the forward interest rate should follow the diffusion process, and that *it* should be the underlying asset for Black's models. But if interest rates follow a diffusion process, then again, negative interest rates could exist.

**4.** The BSOPM assumes that interest rates are constant (or at least nonstochastic). However, this cannot be the case for options on debt instruments, whose prices change because interest rates change. One way around this when one is valuing options on *bonds* is to assume that the short-term interest rate between today and expiration is constant, but the long term rate driving the bond's price is stochastic. This assumption becomes less appealing when one is pricing options on short-term debt instruments.

**5.** The BSOPM assumes that the volatility of the underlying asset is constant. However, the volatility of bond prices is a function of its duration, which is related to its maturity. All else equal, bond price volatilities decline as they near maturity. Estimating the bond's volatility is difficult because it changes both as time passes and as interest rates change. If a spot interest rate is serving as the underlying asset, its volatility will generally increase as time passes (short-term interest rates are more volatile than long-term interest rates). The use of a forward interest rate as the underlying asset keeps constant its duration and volatility.

Despite these shortcomings, however, a variation of the BSOPM developed by Black (1976) is often used to price European-style interest rate options, as well as interest rate caps, interest rate floors, and European options on swaps. However, users of this model are implicitly making one of two assumptions. Either they are assuming that interest rates are constant or they are assuming a lognormal probability distribution for the underlying asset (an interest rate or a bond price).

### 20.2.4.2 Alternative Models Used to Price Interest Rate Derivatives

Given the problems and complications just outlined, those who need accurate values of debt options generally do not use the BSOPM. Instead, numerical methods are chosen. In particular, many use a "lattice" approach. That is, the user builds a "tree" that represents the various evolutionary paths of the interest rate, bond price, or the term structure of interest rates. To help picture this approach, think of how the binomial option pricing model creates a "tree" of prices.[15]

One class of models used in pricing interest rate derivatives is known as the equilibrium class. The models in this class use continuous time mathematics to describe how the "short rate" evolves over time. In these models, the "short rate" is very short. Indeed, it is sometimes called the instantaneous short rate because it is over an infinitely small interval. These models differ from one another in their assumptions of how the short rate evolves over time (i.e., what stochastic process does the short rate follow?).[16] Examples of such models are those of Vasicek (1977), Rendleman and Bartter (1980), and Cox, Ingersoll, and Ross (1985).