CHAPTER 7

From Value-at-Risk to Stress Testing

The measurement of market risk has evolved from simple naïve indicators that distort the measurement of risk, such as the face value or "notional" amount of an individual security, through more complex measures of price sensitivities such as the basis point value or duration approach of a bond (Chapter 6), to sophisticated risk measures such as the latest value-at-risk (VaR) methodology (Figure 7-1) for whole portfolios of securities.

FIGURE 7-1

Traditional Measures of Market Risk

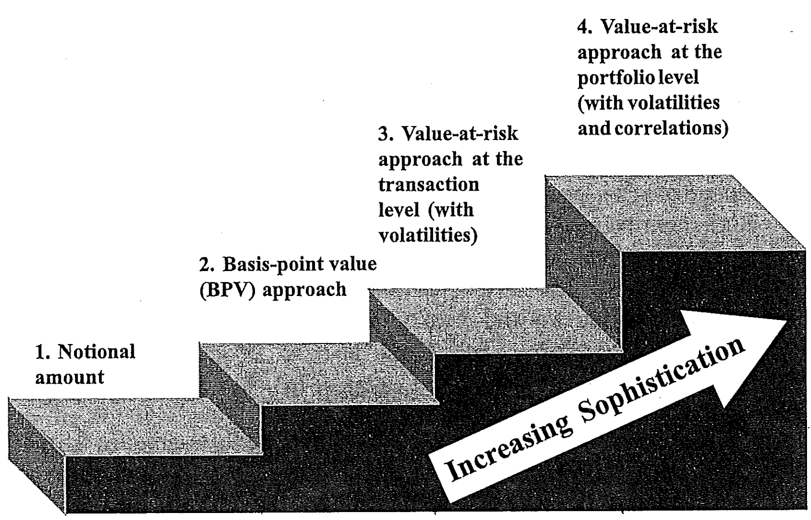

In this chapter we'll explain the principles that lie behind VaR and make clear the strengths and weaknesses of the approach in nonmathematical language. We'll also look at some specialist measures of risk for derivatives (the "Greeks") and at a key supplement to any VaR approach: stress testing and worst-case scenarios.

VaR has proved to be a very powerful way of assessing the overall market risk of trading positions over a short horizon, such as a 10-day period, and under "normal" market conditions. In effect, the methodology allows us to capture in a single number the multiple components of market risk, such as curve risk, basis risk, and volatility risk.

However, each time there is turmoil in the world's markets, the limitations of even the most sophisticated measures of market risk are revealed. VaR has proved unreliable as a measure of risk over long time periods or under abnormal market conditions. The danger posed by exceptional market shocks such as the crisis in the world markets in 1998 that capsized the giant U.S. hedge fund Long-Term Capital Management (LTCM)—shocks that are often accompanied by a drying up of market liquidity—can be captured only by means of supplemental methodologies.

Worries that VaR and other industry-standard risk measures might even exacerbate market volatility have also surfaced. Some influential commentators argue that the herd mentality that is so typical of the financial industry means that market-sensitive risk management systems, such as VaR, actually make markets less stable and more prone to crisis. This is because financial institutions may have to sell assets in the affected classes when markets become volatile in order to keep within the VaR limits set by senior management; this depresses market prices even further and increases the volatility and correlation of the risk factors for these assets. This, in turn, might cause another set of financial institutions to exceed their VaR limits, forcing them to reduce their exposure by selling still more of the same assets—perpetuating a vicious circle.

It's a controversial argument, but then VaR is controversial in so many ways. Let's first take a quick look at the evolution of measures of market risk in the derivative markets before we explore how VaR is calculated, and its resulting strengths and weaknesses.

THE NOTIONAL AMOUNT APPROACH

Until relatively recently, major banks often assessed the amount of market risk generated by their trading desks in terms of the notional or nominal amounts of the portfolio held by the desk. For example, the risk of a portfolio might be assessed with reference to the fact that it contained $30 million of government debt or $30 million of options on the equity of, say, a telecom company. These flawed nominal measures were often routinely presented to senior management and the board as measures of market risk. This is an appealingly simple approach, but it is fatally flawed, since it does not

- Reflect the fact that different assets have vastly different price volatilities (e.g., government bonds are much less likely to fluctuate violently in price than are telecom stocks)
- Take into account the tendency for the value of different assets in the portfolio to rise and fall at the same time (i.e., the "correlation" of the assets in the portfolio)
- Differentiate between short and long positions that might cancel one another out or partially hedge one another (e.g., a long position in a forward contract on the euro with notional value of $100 million maturing in June, and a short position in a forward contract on the euro with a notional value of $50 million maturing in July)

In the case of derivative positions, there are often very large discrepancies between the notional amount, which may be huge, and the true amount of market exposure, which is often small. For example, two call options on the same underlying instrument with the same notional value and the same time to expiration may have very different market values if their strike prices are different—the first option may be deep in-the-money, and the other one may be deep out-of-the money. The first option might be very valuable, while the second might be almost worthless, meaning that they represent very different risk exposures.

As another example, imagine a situation in which interest-rate swaps are written with many different counterparties, and some of these swaps are being used to hedge the market risk exposure created by the other swaps. In this instance, the deals are designed to cancel each other out in terms of their effect on the aggregate market risk in the portfolio. Adding up the notional amounts of the deals will generate an entirely misleading picture of the market risk of the portfolio (although it will offer some indication of overall credit risk exposure).

PRICE SENSITIVITY MEASURES FOR DERIVATIVES

In Chapter 6 we looked at some of the specific measures of market risk in the interest-rate and bond markets. But bond-market traders are not the only practitioners who depend on market-specific risk measures. More recently, practitioners in the derivative markets have developed their own specialized risk measures to describe the sensitivities of derivative instruments to various risk factors. The risk measures are named after letters in the Greek alphabet, and are therefore known collectively as the "Greeks." How do these measures relate to the risk measures that we discussed in Chapter 6?

First, consider a European call option on an individual stock that does not pay any dividend. According to the classic Black-Scholes formula for option pricing, the price of this option is a function of the stock price, the risk-free rate of interest, the instantaneous volatility of the stock return, the strike price, and the option's maturity.

In the option price equation, the stock price plays the same role as the yield in the bond-price relationship that we described in Chapter 6. The sensitivities of the call option price with respect to the stock price are known as the delta and gamma, so we can think of the delta and gamma price risks of a derivative as analogous to the duration and convexity of a bond, respectively. Table 7-1 gives the definitions of the Greeks in more detail.

The list of sensitivities for derivatives in Table 7-1 is longer than a similar list for a standard bond. This is because the value of a derivative is affected by additional risk factors, such as volatility, the discount rate, the passage of time, and, when several risk factors are involved, the correlation between the risk factors.

Weaknesses of the Greek Measures

Traders on options desks use the Greeks to monitor the sensitivities of their market positions and to discuss risk with trading desk risk managers. But each of the sensitivities measured by the Greeks provides only a partial measure of financial risk. The measurements of delta, gamma, and vega complement one another, but they cannot be aggregated to produce an overall measure of the risk generated by a position or a portfolio. In particular,

- Sensitivities cannot be added up across risk types, e.g., the delta and gamma risk of the same position cannot be summed.

TABLE 7-1

The Greek Alphabet for a European Equity Call Option

Delta, or price risk	Delta measures the degree to which an option's value is affected by a small change in the price of the underlying instrument.
Gamma, or convexity risk	Gamma measures the degree to which the option's delta changes as the reference price underlying the option changes. The higher the gamma, the more valuable the option is to its holder. For a high-gamma option, when the underlying price increases, the delta also increases, so that the option appreciates more in value than a gamma-neutral position. Conversely, when the underlying price falls, the delta also falls, and the option loses less in value than if the position were gamma neutral. The reverse is true for short positions in options: high-gamma positions expose their holders to more risk than gamma-neutral positions.
Vega, or volatility risk	Vega measures the sensitivity of the option value to changes in the volatility of the underlying instrument. A higher vega typically increases the value of the option to its holder.
Theta, or time decay risk	Theta measures the time decay of an option. That is, it reflects how much the value of the option changes as the option moves closer to its expiration date. Positive gamma is usually associated with negative time decay, i.e., a natural price attrition of the option as its maturity declines.
Rho, or discount rate risk	Rho measures the change in value of an option in response to a change in interest rate, more specifically, a change in the zero-coupon rate of the same maturity as the option. Typically, the higher the value of rho, the lower the value of the option to its holder.

- Sensitivities cannot be added up across markets, e.g., one cannot sum the delta of a euro/U.S. dollar call and the delta of a call on a stock index.

Since the sensitivities cannot be aggregated, they cannot be used to assess the magnitude of the overall loss that might arise from a change in the risk factors. As a consequence,

- Sensitivities cannot be used directly to measure the amount of capital that the bank is putting at risk.

- Sensitivities do not facilitate financial risk control. Position limits expressed in terms of delta, gamma, and vega are often ineffective, since they do not translate easily into a "maximum dollar loss acceptable" for the position.

This explains the desire for a comprehensive measure of market risk for individual securities and for portfolios. Value-at-risk is one answer to this quest for a consistent measure of market risk.

DEFINING VALUE-AT-RISK

Value-at-risk (VaR) can be defined as the worst loss that might be expected from holding a security or portfolio over a given period of time (say a single day, or 10 days for the purpose of regulatory capital reporting), given a specified level of probability (known as the *confidence level*).

For example, if we say that a position has a daily VaR of $10 million at the 99 percent confidence level, we mean that the realized daily losses from the position will on average be higher than $10 million on only one day in every 100 trading days (i.e., two to three days each year).

This means that VaR is *not* the answer to the simple question: How much can I lose on my portfolio over a given period of time? The answer to this question is "everything," or almost the entire value of the portfolio! Such an answer is not very helpful in practice: it is the correct answer to the wrong question. If all markets collapse at the same time, then naturally prices may plunge and, at least in theory, the value of the portfolio may drop near to zero.

Instead, VaR offers a probability statement about the potential change in the value of a portfolio resulting from a change in market factors over a specified period of time. Crucially, the VaR measure also does not state by *how much* actual losses are likely to exceed the VaR figure; it simply states how likely (or unlikely) it is that the VaR measure will be exceeded.

Most VaR models are designed to measure risk over a short period of time, such as one day, or 10 days in the case of the market-risk measurements required by the regulators for regulatory capital. The confidence level for the calculation of market risk introduced by the Basel Committee in 1998 (BIS 1998) is set at 99 percent. However, for the purposes of allocating internal capital, VaR may be derived at a higher confidence level, say 99.96 percent; this level of confidence is consistent with the level of confidence inherent in an AA credit rating awarded by a public ratings agency.

There are two key steps in calculating VaR. First, derive the forward distribution of the portfolio, or the returns on the portfolio, at the chosen horizon (in this case, one day). We describe later on how this distribution can be derived using three different approaches: historical price distributions (nonparametric VaR); assumptions about normal distributions (parametric VaR); and Monte Carlo analysis.

This distribution is then plotted out as the curve shown in Figure 7-2. This figure shows us how likely it is (vertical axis) that losses of a particular dollar value (horizontal axis) will occur.

Second, identify the required percentile of this distribution so that a particular loss number can be read off. We've selected the first percentile of the distribution in Figure 7-2 because, in this example, we assumed that management has asked for a VaR number measured at the 99 percent confidence level. In Figure 7-2, we also assume that the distribution is a normal bell-shaped curve, rather than a distribution that is skewed toward particularly light or heavy losses. Thus, a confidence level of 99 percent corresponds to a VaR of 2.33 standard deviations.

FIGURE 7-2

Defining Value-at-Risk

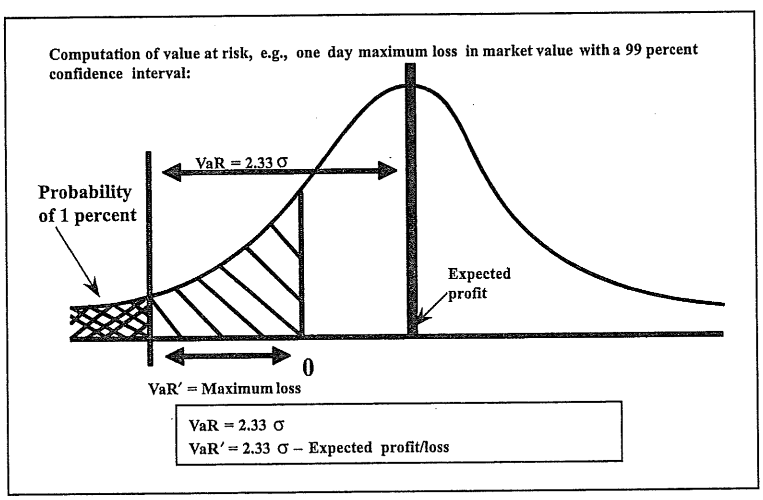

If the confidence level had been set at 99.96 percent, then we would have calculated the 4-basis-point (bp) quantile, and we would have ended up with a larger number for the VaR. (Note that the extent to which the VaR number rises as confidence levels are set more stringently depends upon the shape of the distribution.)

The VaR of the position or portfolio is simply the maximum loss at this 99 percent confidence level, measured relative to the expected value of the portfolio at the target horizon. That is, VaR is the distance of the first percentile from the mean of the distribution.

VaR = expected profit/loss – worst-case loss at the 99 percent confidence level

An alternative and even simpler definition of VaR is that it represents the worst-case loss at the 99 percent confidence level, i.e.,

VaR' = worst-case loss at the 99 percent confidence level

VaR' is also known as *absolute VaR*. However, only our first definition of VaR is consistent with economic capital attribution and the kind of risk-adjusted return on capital (RAROC) calculations we describe in Chapter 15. (Essentially, this is because capital needs to be provided only as a cushion against unexpected losses; in VaR the expected profit or loss is already priced in, and accounted for, in the return calculation.)

So, how exactly does the VaR number relate to economic capital and to regulatory capital? VaR represents the economic capital that shareholders should invest in the firm (or set aside against a particular position or portfolio) to limit the probability of default to a given predetermined level of confidence. Regulatory capital, on the other hand, is the minimum amount of capital imposed by the regulator, as described in Chapter 3. Even when regulatory capital measures are based on a VaR calculation rather than on much simpler rules, economic capital differs from regulatory capital because the confidence level and the time horizon chosen are usually different. For example, when banks are determining their economic capital for market risk, they may choose a higher confidence level than the 99 percent set by the regulator. They may also vary the time horizon when making economic capital calculations, perhaps using one day for very liquid positions, such as a government bond, and as much as several weeks for illiquid positions, such as long-dated over-the-counter equity derivatives. By contrast, the regulator arbitrarily sets the time horizon at 10 days for any position in the trading book.

From 1-Day VaR to 10-Day VaR

VaR is often used to manage market risk over a 1-day time horizon. For this purpose, it's necessary to derive VaR from the daily distribution of the portfolio values. However, we mentioned earlier that the regulators have set a time horizon of 10 days for the purpose of VaR calculations that are used to report regulatory capital requirements. Ideally, this "10-day VaR" would be derived from a corresponding distribution of results over a 10-day horizon. This is problematic, however, as it implies that the time series of data used for the analysis must be much longer—indeed, 10 times longer—than that employed in any one-day VaR analysis. As a result, many banks employ a work-around that allows them to derive an approximation of 10-day VaR from daily VaR data by multiplying the daily VaR by the square root of time (here, 10 days). The "square root of time" rule is endorsed by the regulators; it should be noted, however, that it is not a sound practical approach and remains something of a rule of thumb.

HOW IS VaR USED TO LIMIT RISK IN PRACTICE?

VaR is an aggregate measure of risk across all risk factors. A special attraction is that it can be calculated at each level of activity in the business hierarchy of a company. For example, it can be calculated for each activity (e.g., trading desk) at both the business unit level (e.g., equity trading) and the level of the firm as a whole.

At the level of the firm, VaR offers a good way of representing the (short-term) "risk appetite" of the firm, since it measures the maximum loss that the firm might incur, under normal market conditions, over a short period of time (in effect, 1 to 10 days). The risk appetite of the firm over a longer period of time, say a quarter, is usually set in terms of a worst-case scenario analysis. For example, the board of the bank can set a limit on the maximum loss that it is prepared to tolerate over a quarter if the worst-case crisis that the bank's risk managers think plausible in that period should, in fact, occur.

At many financial institutions, the board of directors sets an overall VaR limit whose control is delegated to the chief executive officer (CEO). In practice, this control is often delegated, in turn, to a risk management committee chaired by the CEO. In many banks, the risk management committee appoints a chief risk officer (CRO) or similar risk executive to report on firmwide risk and therefore help maintain effective control of this limit. We discuss this cascade of accountability in more detail in Chapter 4. Box 7-1 reviews the strengths of VaR, not only as a measurement tool, but also as a managerial instrument.

BOX 7-1

VaR IS FOR MANAGING, AS WELL AS MEASURING, RISK

In the main text, we highlight the problems inherent in the simplifying assumptions that must be made whenever a VaR number is calculated. In this box, let's remind ourselves of the great strengths of VaR and its wide range of uses:

- *VaR provides a common, consistent, and integrated measure of risk across risk factors, instruments, and asset classes.* For example, it allows managers to measure the risk of a fixed-income position in a way that is comparable and consistent with their risk assessment of an equity derivative position. VaR also takes into account the correlations between the various risk factors, somewhat in the spirit of portfolio theory.
- *VaR can provide an aggregate measure of risk and risk-adjusted performance.* This single number can then be easily translated into a capital requirement. VaR can also be used to reward employees on the basis of the risk-adjusted return on capital generated by their activities. In other words, it can be used to measure risk-adjusted performance (see Chapter 15).
- *Business-line risk limits can be set in terms of VaR.* These limits can be used to ensure that individuals do not take more risk than they are allowed to take. Risk limits expressed in VaR units can easily be aggregated up through the firm, from the business line at trading desk level to the very top of the corporation. The drill-down capability of a VaR system allows risk managers to detect which unit is taking the most risk, and also to identify the type of risk to which the whole bank is most exposed, e.g., equity, interest-rate, currency, or equity vega.
- *VaR provides senior management, the board of directors, and regulators with a risk measure that they can understand.* Managers and shareholders, as well as regulators, can decide whether they feel comfortable with the level of risk taken on by the bank in terms of VaR units. VaR also provides a framework for assessing, *ex ante*, investments and projects on the basis of their expected risk-adjusted return on capital.

(continued on following page)

BOX 7-1 *(Continued)*

- *A VaR system allows a firm to assess the benefits from portfolio diversification within a line of activity and across businesses.* VaR allows managers to assess the daily revenue volatility they might expect from any given trading area, but it also allows them to compare the volatilities of different business areas, such as equity and fixed-income businesses, so that they can understand better how each business line offsets, or contributes to, the revenue volatility of the whole firm.
- *VaR has become an industry-standard internal and external reporting tool.* VaR reports are produced daily for managers of business lines, and are then aggregated for senior management. VaR is also communicated to the regulators and has become the basis for calculating regulatory capital in some areas of risk measurement. The rating agencies take VaR calculations into account in establishing their ratings of banks. Increasingly, VaR and the results of the back testing of VaR are published in banks' annual reports as a key indicator of risk.

Figures 7-3*A* and *B* helps us to visualize what VaR measures mean in practice, and how they are used to manage risk on a trading desk. (For this example, we'll stick with the nonparametric, or historical VaR, approach to calculating VaR, one of a set of calculation approaches that we explain in more detail later on.) In this illustrative example, the average daily revenue for our example bank's trading portfolio during 1998 was C$0.451 million. But what we are really interested in is the distribution of the bank's gains and losses, represented in Figure 7-3*B*, which tells us how frequently the bank incurred each loss amount. The first percentile of the historical distribution represented in Figure 7-3*B*, i.e., the cutoff point on this distribution such that only 1 percent of the daily revenues lies on the distribution's left-hand side, is C$25.919 million. This represents VAR', or the absolute VaR to a 99 percent level of confidence. From our earlier discussion, we know that to work out the true one-day historical VaR for the portfolio, we need to take the expected profit or loss into account. So our VaR number to a 99 percent level of confidence based on this 12-month set of data is $0.451 - (-25.919) =$ C$26.37 million.

Now let's turn again to Figure 7-3*A* to discuss how VaR limits are used as a practical tool for managing market risk. Retrospectively, 1998 is an interesting year for the purpose of analyzing risk management decisions. Market participants were surprised by the severe market disruptions in August 1998 after the Russian government defaulted on its debt.

Liquidity suddenly evaporated from many financial markets, causing asset prices to plunge and producing large losses for many financial institutions (and the near-collapse of the U.S. hedge fund LTCM). Figure 7-3*A* shows the aggregate VaR creeping up slowly during the first part of the year, then increasing substantially during May and June as market volatility edged higher. During that period, the Senior Risk Committee limit for the bank remained at $58 million, well above the daily VaR (Figure 7-3*A*). As risk kept increasing during the summer, the Senior Risk Committee limit was lowered to $38 million in July before the August market crisis. At the peak of the market disruption during the month of August, the new VaR limit became binding, putting pressure on the bank's trading businesses to lower their risks. We can see from the figure that the bank experienced substantial trading losses during the month of August, and after this the VaR limit was further reduced in order to oblige the trading businesses to reduce their risk exposure still further.

As a general point, VaR limits for individual business lines such as trading desks must be set at a level consistent with the firm's overall VaR limit. Otherwise the risk exposures of all the business units might remain

FIGURE 7-3A

Net Daily Trading Revenues during 1998 versus One-Day VaR at the 99 Percent Confidence Level.

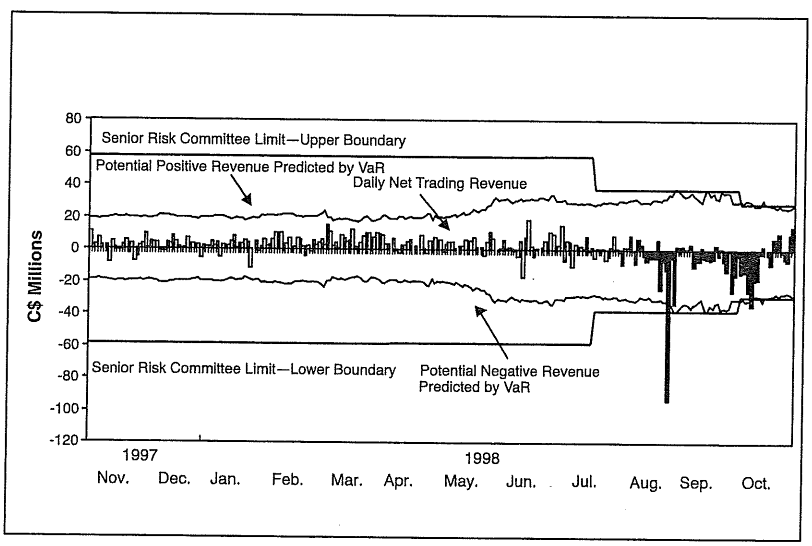

FIGURE 7-3B

Net Daily Trading Revenues for 1998 (C$ Millions)

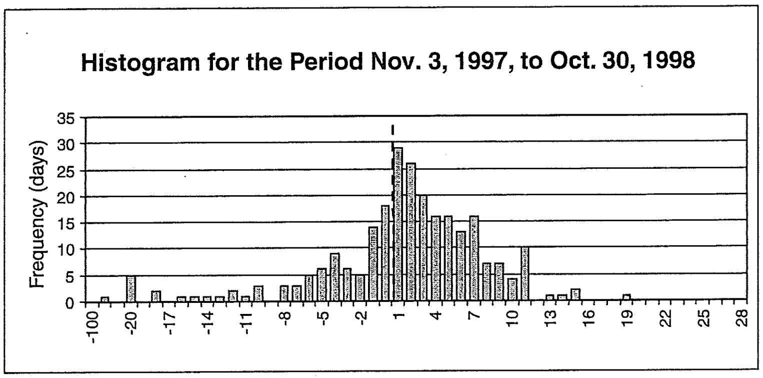

within their own limits, while the firm's aggregate risk breached the overall VaR limit set at the top of the firm.

HOW DO WE GENERATE DISTRIBUTIONS FOR CALCULATING VAR?

To calculate VaR, we first need to select the factors that drive the volatility of returns in the trading or investment portfolio. We can then use these risk factors to generate the forward distribution of the portfolio values at the risk horizon (or, equivalently, the distribution of the changes in the value of the portfolio). Only after generating the distribution can we calculate the mean and the quantiles of this distribution to arrive at the portfolio VaR.

1. Selection of the Risk Factors

The change in the value of the portfolio is driven by changes in the market factors that influence the price of each instrument. The relevant risk factors depend on the composition of the portfolio. The selection of risk factors is straightforward for a simple security, but it requires judgment for more complex products.

In the case of a simple security, such as a U.S. $/euro forward, the value of the position is affected only by the U.S. $/euro forward rate. In the case of a U.S. $/euro call option, the value of the position depends not

TABLE 7-2

Example of a Selection of Risk Factors

U.S.$/euro forward	U.S.$/euro option
■ U.S.$/euro forward rate	■ US$/euro exchange rate ■ US$ interest rates ■ Euro interest rates ■ US$/euro volatility

only on the U.S. $/euro exchange rate, but also on the dollar and euro interest rates over the maturity of the option and on the U.S. $/euro volatility (Table 7-2).

In the case of a stock portfolio, the risk factors are the prices of the individual stocks that make up the portfolio. For a bond portfolio, the choice of the risk factors depends on the degree of "granularity" that one needs in order to understand the risk in hand. For example, the risk factor for each bond might simply be its yield to maturity, as described in Chapter 6. Alternatively, it might be a selection of zero-coupon rates on the risk-free term structure of interest rates for each currency. The selection might comprise the overnight, 1-month, 3-month, 6-month, 1-year, 3-year, 5-year, 10-year, and 30-year zero-coupon rates, as well as the spread in prices between different issuers for the same terms (so that the calculation captures issuer risk).

2. Choice of a Methodology for Modeling Changes in Market Risk Factors

Having identified the risk factors that generate the volatility in the portfolio's returns, the risk analyst must choose an appropriate methodology for deriving the distribution. There are three alternatives:

- The analytic variance-covariance approach
- The historical simulation approach
- The Monte Carlo simulation approach

Analytic Variance-Covariance Approach: Case of a Portfolio Linear in Risks

To simplify the derivation of VaR, we can choose to make certain assumptions. Under the analytic variance-covariance approach or "delta normal" approach, we assume that the risk factors and the portfolio values are log-normally distributed or, equivalently, that their log returns (the log of the returns) are normally distributed. This makes the calculation much simpler, since the normal distribution is completely characterized by its first two moments, and the analyst can derive the mean and the variance of the portfolio return distribution from

- The multivariate distribution of the risk factors
- The composition of the portfolio

A simple example should help make the process clear. Suppose our example portfolio is composed of two stocks, Microsoft and Exxon. In this example, the risk factors that generate the returns in the portfolio are straightforward to identify: the stock prices for each of the companies, the volatility of both stocks, and the correlation coefficient that describes the extent to which the stock prices of Microsoft and Exxon go up and down together.

From historical data on the behavior of the two stocks, we can estimate the simple historical mean and standard deviation of the daily returns for each of the two stocks for each day over a one-year trading period. We could obtain this stock price information from any of the major market information providers, such as Reuters or Bloomberg.

The historical data also allow us to estimate a correlation risk factor for the price relationship between the two stocks. The correlation risk factor is quite important: when the stocks are perfectly correlated, the VaR will be the sum of the VaRs of the individual stocks. Most stocks are not strongly correlated, however, so the VaR tends to be considerably less than the sum of the VaRs of the individual stocks.

Under this approach, remember that we assume that the rates of return on the stocks follow a multivariate normal distribution. This assumption means that we can apply our risk-factor analysis to the present portfolio to generate a distribution of returns of the portfolio into the future. Of course, we must take into account the present price of the portfolio and the percentage of each stock that the portfolio contains.

Having generated the distribution using our five risk factors, we can plot the distribution so that it looks rather like the curve in Figure 7-2, referred to earlier in our discussion. It is now a simple enough matter to read off the VaR number that is relevant to our selected confidence level (e.g., 99 percent), as we described earlier for Figure 7-2.

Our discussion of this approach to calculating VaR begs a major question: how dangerous is our simplifying assumption that returns are normally distributed? In fact, there is a large amount of evidence that many individual return distributions are not normally distributed, but rather exhibit what are known as "fat tails." The term *fat tails* arises out of the shape of certain distributions when plotted on a graph. In these distributions, there are more observations far away from the mean than is the case in a normal or bell-shaped distribution. So whereas a normal distribution tails off quickly (to reflect the rarity of unlikely events), the tail of a fat-tailed distribution remains relatively thick. We can see the difference in Figure 7-4, where the dotted line represents a normal distribution and the continuous line a fatter-tailed distribution.

Fat tails in distributions should worry risk managers because they imply that extraordinary losses occur more frequently than a normal distribution would lead us to believe.

We would expect the VaR derived from a fat-tailed distribution to be higher than that derived from a normal distribution—perhaps much higher. It follows that if we assume that a distribution is normal in our VaR calculation, when in fact it has a fat tail, we are likely to underestimate the VaR number associated with the financial portfolio.

FIGURE 7-4

Comparison of the Normal and a Fat-Tailed Distribution

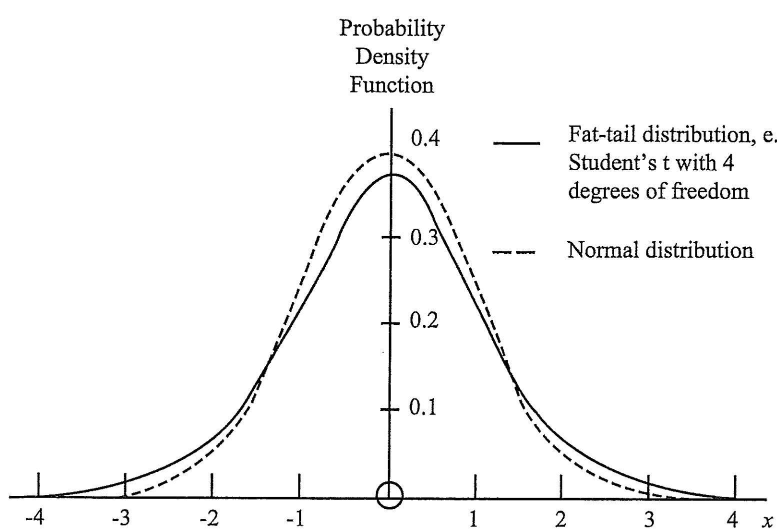

Luckily, even if the returns of an individual risk factor do not follow a normal distribution, we can reasonably expect that the returns of a well-diversified portfolio (i.e., a portfolio subject to many different risk factors) will still exhibit a normal distribution. This effect is explained by the central limit theorem, which tells us that the independent random variables of a well-behaved distribution will have a mean that converges, in large samples, to a normal distribution.

In practice, this result implies that a risk manager can assume that a portfolio has a normal distribution of returns, provided that the portfolio is fairly well diversified and the risk-factor returns are sufficiently independent from one another (even when they are not, themselves, normally distributed).

However, the potential effect on VaR calculations of fat tails, lumpy portfolios, and correlated risk-factor returns should send a warning signal to support staff and senior managers who use VaR numbers to gain comfort about risk levels.

Historical Simulation Approach

The historical simulation approach to VaR calculation is conceptually simple and does not oblige the user to make any analytic assumptions about the distributions. However, at least two or three years of historical data are necessary to produce meaningful results. We've already applied the principles of this approach in our earlier example of the VaR number associated with trading revenues in 1998. In this special case, there was only one risk factor: the daily trading revenue of the firm. In the following, we consider the more usual case: analyzing the VaR of a whole portfolio of securities.

First, the changes in relevant market prices and rates (the risk factors) that have been seen are analyzed over a specified historical period, say, two years. The portfolio under examination is then revalued, using changes in the risk factors derived from the historical data, to create the distribution of the portfolio returns from which the VaR of the portfolio can be derived. Each daily simulated change in the value of the portfolio is considered as an observation in the distribution.

Three steps are involved:

- Select a sample of actual daily risk factor changes over a given period of time, say 500 days (i.e., two years' worth of trading days), using the same period of time for all the factors.

- Apply those daily changes to the current value of the risk factors, revaluing the current portfolio as many times as the number of days in the historical sample. Sum these changes across all positions, keeping the days synchronized.
- Construct the histogram of portfolio values and identify the VaR that isolates the first percentile of the distribution in the left-hand tail (assuming VaR is derived at the 99 percent confidence level).

Let's illustrate the approach using an example. Assume that the current portfolio is composed of a three-month U.S. \$/DM call option. The market risk factors for this position are

- U.S. \$/DM exchange rate
- U.S. \$ three-month interest rate
- DM three-month interest rate
- Three-month implied volatility of the U.S. \$/DM exchange rate

In the following, we neglect the impact of interest-rate risk factors and consider only the level of the exchange rate and its volatility. The first step is to report daily observations of the risk factors we've selected over the past 100 days, as shown in abbreviated form in columns 2 and 3 of Table 7-3.

Historical simulation, like Monte Carlo simulation, requires the repricing of the position in question using the historical distribution of the risk factors. In this example, we use the Black-Scholes model adapted by

TABLE 7-3

Historical Market Values for the Risk Factors Over the Last 100 Days

Day (t)	U.S.\$/DM (FX_t)	FX Volatility (σ_t)
−100	1.3970	0.149
−99	1.3960	0.149
−98	1.3973	0.151
...	...	...
−2	1.4015	0.163
−1	1.4024	0.164

TABLE 7-4

Simulating Portfolio Values Using Historical Data (Current Value of the Portfolio: $1.80)

	Change from Current Value ($1.80)
Alternative price 100 = $C(FX_{100}; \sigma_{100})$ = $1.75	−$0.05
Alternative price 99 = $C(FX_{99}; \sigma_{99})$ = $1.73	−$0.07
Alternative price 98 = $C(FX_{98}; \sigma_{98})$ = $1.69	−$0.11
....	
Alternative price 2 = $C(FX_{2}; \sigma_{2})$ = $1.87	+$0.07
Alternative price 1 = $C(FX_{1}; \sigma_{1})$ = $1.88	+$0.08

TABLE 7-5

Identifying the First Percentile of the Historical Distribution of the Portfolio Return

Rank	Change from Current Value
100	−$0.05
99	−$0.07
98	−$0.11
...	...
2	+$0.07
1	+$0.08

Garman-Kholhagen (1983) to currency options. The results of this step are reported in Table 7-4.[1]

The last step consists of constructing the histogram of the portfolio returns based on the last 100 days of history or, equivalently, sorting the changes in portfolio values to identify the first percentile of the distribution. Table 7-5 shows the ranking of the changes in the value of the portfolio. Using this, we identify the first percentile as -$0.07.

1. F. Black, and M. Scholes, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* 81, 1973, pp. 637-654; M. B. Garman and S. Kolhagen, "Foreign Currency Option Values," *Journal of International Money and Finance* 2, December 1983, pp. 231-237.

FIGURE 7-5

VaR from Historical Simulations

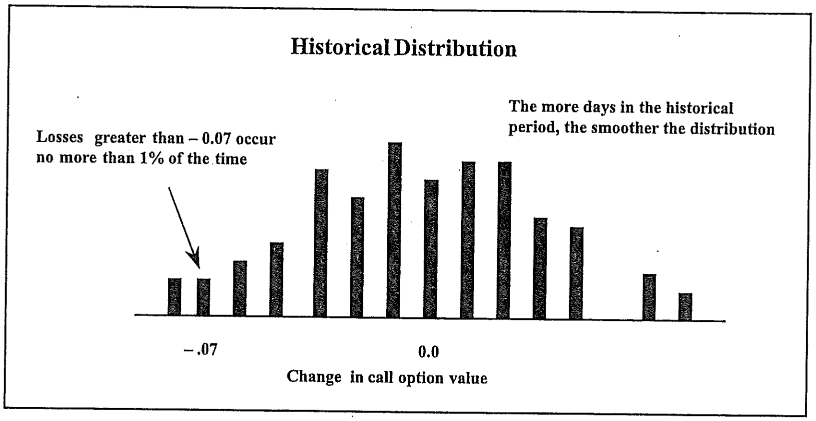

Figure 7-5 shows the histogram of these values. VaR (1; 99) at the 99 percent confidence level is simply the distance to the mean ($0.01) of the first percentile, i.e., VaR (1; 99) = $0.08, while absolute VaR is the first percentile itself, i.e., VaR' (1; 99) = $0.07. Note that this histogram is similar to the histogram that we derived for daily trading revenues in Figure 7-3*b*.

This three-step procedure can easily be generalized to any portfolio of securities.

The major attraction of historical simulation is that the method is completely nonparametric (i.e., we don't need to worry about setting parameters) and does not depend on any assumptions about the distribution of the risk factors. In particular, we do not need to assume that the returns of the risk factors are normally distributed and independent over time.

The nonparametric nature of historical simulation also obviates the need to estimate volatilities and correlations. Historical volatilities and correlations are already reflected in the data set, so all we need to calculate are the synchronous risk-factor returns over a given historical period. Historical simulation has also no problem accommodating fat tails in distributions, since the historical returns already reflect actual synchronous moves in the market across all risk factors. Another advantage of historical simulation over the variance-covariance approach is that it allows the analyst to calculate confidence intervals for VaR.

The main drawback of historical simulation is its complete dependence on a particular set of historical data, and thus on the idiosyncrasies of this data set. The underlying assumption is that the past, as captured in this historical data set, is a reliable representation of the future (i.e., the past is prologue). This implicitly presumes that the market events embedded in the historical data set will be reproduced in the months to come. However, the historical period may cover events, such as a market crash or, conversely, a period of exceptionally low price volatility, that are unlikely to be repeated in the future. Historical simulation may also lead to a distorted assessment of the risk if we employ the technique regardless of any structural changes anticipated in the market, such as the introduction of the euro in the foreign exchange markets at the beginning of 1999.

Another practical limitation of historical simulation is data availability. One year of data corresponds to only 250 data points (trading days) on average, i.e., 250 scenarios. By contrast, Monte Carlo simulations usually involve at least 10,000 simulations (i.e., scenarios). Employing small samples of historical data inevitably leaves gaps in the distributions of the risk factors and tends to underrepresent the tails of the distributions, i.e., the occurrence of unlikely but extreme events.

Monte Carlo Approach

Monte Carlo simulation consists of repeatedly simulating the random processes that govern market prices and rates. Each simulation (scenario) generates a possible value for the portfolio at the target horizon (e.g., 10 days). If we generate enough of these scenarios, the simulated distribution of the portfolio's values will converge toward the true, although unknown, distribution. The VaR can be easily inferred from the distribution, as we described earlier.

Monte Carlo simulation involves three steps:

1. *Specify all the relevant risk factors*. As in the other approaches, we need to select all the relevant risk factors. In addition, we have to specify the dynamics of these factors, i.e., their stochastic processes, and we need to estimate their parameters (volatilities, correlations, mean reversion factors for interest-rate processes, and so on).
2. *Construct price paths*. Price paths are constructed using random numbers produced by a random-number generator. For a simple portfolio without complex exotic options, the forward distribu-

tion of portfolio returns at a 10-day horizon can be generated in one step. Alternatively, if the simulation is performed on a daily basis, a random distribution is drawn for each day to calculate the 10-day cumulative impact.

When several correlated risk factors are involved, we need to simulate multivariate distributions. Only in the case where the distributions are independent can the randomization be performed independently for each variable.

3. *Value the portfolio for each path (scenario).* Each path generates a set of values for the risk factors for each security in the portfolio that are used as inputs into the pricing models. The process is repeated a large number of times, say 10,000 times, to generate the distribution, at the risk horizon, of the portfolio return. This step is equivalent to the corresponding procedure for historical simulation, except that Monte Carlo simulation can generate many more scenarios than historical simulation.

VaR at the 99 percent confidence level is then simply derived as the distance to the mean of the first percentile of the distribution, as for our other calculation methods.

Monte Carlo simulation is a powerful and flexible approach to VaR. It can accommodate any distribution of risk factors to allow for fat-tailed distributions, where extreme events are expected to occur more commonly than in normal distributions, and "jumps" or discontinuities in price processes. For example, a process can be described as a mixture of two normal distributions or as a jump-diffusion process where the number of jumps in any time interval is governed by, say, a Poisson process (both processes are consistent with fat tails).

Monte Carlo simulation, like historical simulation, allows the analyst to calculate the confidence interval of VaR, i.e., the range of likely values that VaR might take if we repeated the simulation many times. The narrower this confidence interval, the more precise the estimate of VaR. Monte Carlo simulation has a particular advantage here: it is easy to carry out sensitivity analyses by changing the market parameters used in the analysis, such as the term structure of interest rates.

One disadvantage of the Monte Carlo approach is that the analyst must be able to estimate the parameters of the distributions, such as the means, the variances, and the covariances. The major limitation of the approach, however, is more pragmatic: the amount of computer resources it consumes. Variance-reduction techniques can be used to reduce the computational time, but Monte Carlo simulation remains very computer intensive and cannot be used to calculate the VaR of very large and complex portfolios.

Pros and Cons of the Different Approaches

Each of the approaches we have described has advantages and limitations; no single technique is "perfect," and no technique dominates the others.

For this reason, it is important that financial professionals and managers who rely on VaR numbers to measure risk—or to gain comfort about the risks that an institution is taking—be familiar with the basic principles of the VaR calculation. Increasingly, equity analysts and investors also need to understand these numbers if they are to assess the information that a bank makes public about its risk profile.

Table 7-6*A* through *C* summarizes the pros and cons of the different approaches. Together with the information contained in this chapter, it can be used to frame questions about how any particular VaR number has been produced.

Above all, people using VaR numbers must remember that they are not a "magic bullet" for measuring and managing risk. In the right hands, VaR techniques help to offer risk analysts a rational and comparable snapshot of the risk of a particular position or portfolio. But like every financial measure, VaR numbers in the wrong hands can be used to mislead and obfuscate. Their reliability as a decision-making tool depends upon the skill and experience of the analyst, the nature of the problem that is being explored, and the ability of decision makers to ask intelligent questions about meaning and provenance.

TABLE 7-6A

Pros and Cons of the Variance-Covariance Approach

Pros	Cons
Computationally efficient; it takes only a few minutes to run the position of the entire bank.	Assumes normality of the return portfolio.
Because of central limit theorem, the methodology can be applied even if the risk factors are not normal, provided the factors are numerous and relatively independent.	Assumes that the risk factors follow a multivariate log-normal distribution, and thus does not cope very well with "fat-tailed" distributions.

(continued on following page)

TABLE 7-6A *(Continued)*

Pros and Cons of the Variance-Covariance Approach

Pros	Cons
No pricing model is required; only the Greeks are necessary, and these can be provided directly by most of the systems that already exist within banks (i.e., the legacy systems).	Requires the estimation of the volatilities of the risk factors as well as the correlations of their returns.
It is easy to handle incremental VaR.	Security returns can be approximated by means of a Taylor expansion. In some instances, however, a second-order expansion may not be sufficient to capture option risk (especially in the case of exotic options).
	Cannot be used to conduct sensitivity analysis.
	Cannot be used to derive the confidence interval for VaR.

TABLE 7-6B

Pros and Cons of the Historical Simulation Approach

Pros	Cons
No need to make any assumption about the distribution of the risk factors.	Complete dependence on a particular historical data set and its idiosyncrasies. That is, extreme events such as market crashes either lie outside the data set and are ignored, or lie within the data set and (for some purposes) act to distort it.
No need to estimate volatilities and correlations; they are implicitly captured by the actual (synchronous) daily realizations of the market factors.	Cannot accommodate changes in the market structure, such as the introduction of the euro in January 1999.
Fat tails of distributions, and other extreme events, are captured so long as they are contained in the data set.	Short data set may lead to biased and imprecise estimation of VaR.
Aggregation across markets is straightforward.	Cannot be used to conduct sensitivity analyses.
Allows the calculation of confidence intervals for VaR.	Not always computationally efficient when the portfolio contains complex securities.

TABLE 7-6C

Pros and Cons of the Monte Carlo Simulation Approach

Pros	Cons
Can accommodate any distribution of risk factors.	Outliers are not incorporated into the distribution.
Can be used to model any complex portfolio.	Computer intensive.
Allows the calculation of confidence intervals for VaR.	
Allows the user to perform sensitivity analyses and stress testing.	

Perhaps the key "lesson learned" about VaR numbers, however, is that they must be supplemented by the methodologies we turn to in the next section: stress testing and worst-case scenarios.

STRESS TESTING AND SCENARIO ANALYSIS

As we discussed earlier, VaR is far from being a perfect measure of risk. Its use and reliability are often dictated by the availability of data; for instance, reliable data for implied volatilities can be obtained only for short maturities. And to facilitate the implementation of a VaR model, especially in the case of the analytic variance-covariance and Monte Carlo approaches, it is common to assume that market conditions will remain normal. Prices and values are assumed to have a "smooth" behavior that excludes the possibility of jumps and other extreme events.

We don't yet know how to construct a VaR model that would combine, in a meaningful way, periods of normal market conditions with periods of market crises characterized by large price changes, high volatility, and a breakdown in the correlations among the risk factors.

Another problem is that VaR is usually calculated within a static framework and is therefore appropriate only for relatively short time horizons, which in turn means that we can't include dynamic liquidity risks in the VaR analysis.

All this means that stress testing and scenario analysis must be used as supplementary methodologies to help us analyze the possible effects of extreme events that lie outside normal market conditions. Regulators view stress testing and scenario analysis as a necessary complement to the use of internal VaR models.

The purpose of stress testing and scenario analysis is to determine the size (though not the frequency) of potential losses related to specific scenarios. The selection of an appropriate scenario is largely based on expert judgment. The scenario may consist of extreme changes in the value of a risk factor (interest rate, exchange rate, equity price, or commodity price), such as a shift of 100 bp in the level of interest rates over the period of a month; replication scenarios that attempt to reproduce the multi-risk-factor effects of extreme historical events; and hypothetical one-time events that depend on imagined future developments.

Risk-Factor Stress Testing

Back in 1995, the Derivative Policy Group recommended some specific guidelines for risk-factor stress testing that help give us a flavor of the range of stresses banks now use to test out their derivative exposures:[2]

- Parallel yield-curve shift of plus or minus 100 bp
- Yield-curve twist of plus or minus 25 bp
- Equity index values change of plus or minus 10 percent
- Currency changes of plus or minus 6 percent
- Volatility changes of plus or minus 20 percent

These extreme price and rate variations may dramatically affect the value of a portfolio with strong nonlinearity and large negative gammas. Portfolios of this kind incur losses whether prices fall or rise, and the magnitude of the losses accelerates as the change in price increases. (Recall that gamma is the nickname for the second derivative of the instrument's value with respect to the value of the underlying asset.)

The regulators also now require that financial institutions run scenarios that capture the specific characteristics of their portfolios, i.e., scenarios that involve the risk factors to which their portfolios are most sensitive. Following the market crisis of the summer of 1998, when the disappearance of liquidity in some financial markets led to several well-publicized losses (we recount the story of the near-collapse of LTCM in Chapter 14), regulators required financial institutions to include liquidity

2. Derivatives Policy Group. A framework for voluntary Oversight, 1995.

risk in their scenario analyses. Many institutions now also apply stress-testing approaches to their credit and operational risk exposures.

Stress-Testing Envelopes

Stress testing is becoming more and more sophisticated, and one of the challenges is to work out a rigorous way of applying different kinds of stress to portfolios in a consistent manner.

Here we'll consider a "stress-envelope" methodology, which combines stress categories with the worst possible "stress shocks" across all possible markets for every business. For example, the methodology might designate seven stress categories corresponding to the various risk categories: interest rates, foreign exchange rates, equity prices, commodity prices, credit spreads, swap spreads, and vega (volatility). For each stress category, the worst possible stress shocks that might realistically occur in the market are defined. In the case of interest rates, for example, the methodology defines six stress shocks to accommodate both changes in the *level* of rates and changes in the *shape* of the yield curve. In the case of credit spreads and equities, there is only one stress shock, i.e., the widening of credit spreads and the fall of equity prices, respectively. All other stress categories make use of one or two stress shocks (increases or decreases in spreads or prices), as shown in Table 7-7.

The number of markets, currencies, and businesses must be determined empirically for each institution.

We can think of the stress envelope itself as the change in the market value of a business position in a particular currency or market in

TABLE 7-7

Stress Categories and the Number of Stress Shocks

	Stress Category	Stress Shocks
1	Interest rates	6
2	Foreign exchange	2
3	Equity	1
4	Commodity	2
5	Credit spreads	1
6	Swap spreads	2
7	Vega (volatility)	2

FIGURE 7-6

The Seven Major Components of the Stress-Envelope Approach

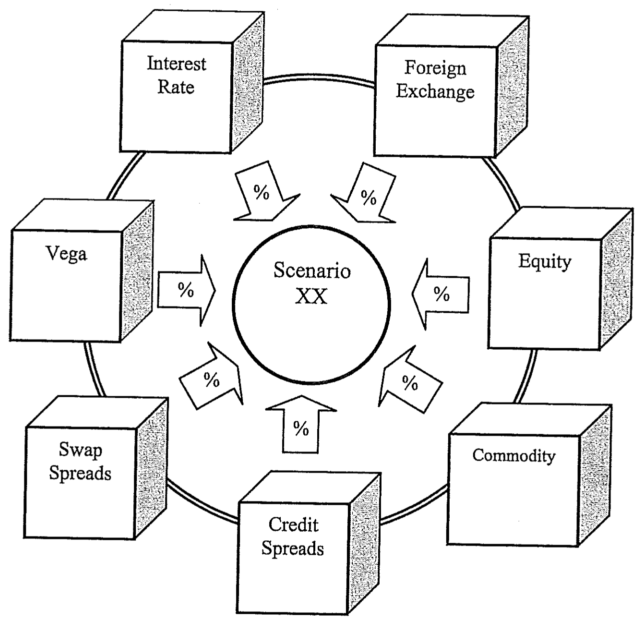

response to a particular stress shock. A scenario can then be created using a combination of several stress shocks (Figure 7-6).

As well as testing hypothetical extreme changes in risk factors, a scenario might also correspond to extreme historical events, such as the stock market crash of October 1987. On the October 19, 1987, stock prices in the United States fell by 23 percent, or approximately 22 times their daily standard deviation. Other historical scenarios might include the failure of the European exchange-rate mechanism in September 1992 or the tightening of monetary policy by the Federal Reserve in the United States in May 1994 (and the subsequent fall in bond prices).

Boxes 7-2 and 7-3 present examples of historical replication scenarios. We might also build hypothetical extreme scenarios to explore the risks of plausible structural changes in the market that are likely to damage a portfolio. For example, we might set values around an event in which the Chinese devalue their currency, causing a crisis in Asia.

BOX 7-2

REPLICATION SCENARIO 1: STOCK MARKET CRASH

As an example of a historical replication scenario, consider a stock market crash reminiscent of the crisis in the global financial markets in October 1987, characterized by a combination of the following events:

- Equity markets around the globe fall by 20 percent on average, with Asian markets, such as Hong Kong, declining by 30 percent, and there is an upward shift in implied volatilities from 20 to 50 percent.
- The U.S. dollar rallies against other currencies as a consequence of a flight to quality. Asian currencies lose up to 10 percent against the dollar.
- Interest rates fall in Western markets. Hong Kong interest rates rise by 40 bp at the long end of the term structure and by 100 bp at the short end.
- Commodity prices drop as a result of fears of a recession; copper and oil prices decline by 5 percent.

BOX 7-3

REPLICATION SCENARIO 2: U.S. MONETARY TIGHTENING

In this example of a historical replication scenario, consider a U.S. inflation scare and a tightening of monetary policy by the U.S. Federal Reserve along the lines of that seen in May 1994, characterized by

- A 100-bp increase in the overnight interest rate and a 50-bp upward shift in the long end of the curve.
- Interest rates also increase in other G-7 countries and Switzerland, but not as much as in the United States.
- G-7 currencies depreciate against the U.S. dollar as investors chase higher rates.
- Credit spreads widen.
- Equity markets decline from 3 to 6 percent, with an upward shift in implied volatilities.

Advantages of Stress Testing and Scenario Analysis

The major benefit of stress testing and scenario analysis is that they show us how vulnerable a portfolio might be to a variety of extreme events. During a market crisis, historical correlations change as volatilities increase. Correlations may suddenly increase dramatically and become +1 as many markets collapse at the same time and liquidity dries up. Alternatively, correlations may approach –1 as markets or instruments move in opposite directions. For example, a market event may trigger a flight to quality, with liquid and illiquid markets exhibiting almost perfect negative correlation.

Each portfolio has specific characteristics that make it vulnerable to a particular scenario and/or stress tests. Obviously, a high-yield bond portfolio is vulnerable to a widening of credit spreads. An equity portfolio that is diversified across many countries and sectors of activities is sensitive to a change in the correlation structure of the world's equity markets. An equity derivative book that is short gamma is vulnerable to a sharp increase in volatility. Stress testing and scenario analyses are very useful in highlighting these unique vulnerabilities for senior management.

Limitations of Stress Testing and Scenario Analysis

Stress testing and scenario analysis are important building blocks in any risk management methodology, but they can't tell us how likely it is that a particular event will come to pass. They also have many other limitations that we must bear in mind:

- Scenarios are based on an arbitrary combination of stress shocks. Yet many such combinations are inconsistent with the basic laws of economics. They may violate, for example, no-arbitrage conditions such as interest-rate parity. When constructing a scenario, it is important to examine the chain of events to make sure that it makes economic sense. The chain of events that may logically follow an initial major shock depends on the context, and may be quite different from one crisis to another. For example, the Asian crisis of the summer of 1997 was quite different from the Asian crisis of the summer of 1998 (triggered by the partial default of Russia).

- The potential number of combinations of basic stress shocks is overwhelming. In practice, only a relatively small number of scenarios can be routinely analyzed. This means that the scenarios have to be selected according to the vulnerabilities of the particular portfolio. Again, the choice is necessarily somewhat arbitrary. The usefulness and accuracy of the diagnosis that emerges from the scenario analysis depends on the judgment and experience of the analysts who design and run these scenarios. Even the best analysts rely on the past as a guide to the future. Yet history is unlikely to repeat itself exactly.
- Market crises unfold over a period of time, during which liquidity may dry up. Yet most scenario analyses are static in nature; i.e., they are one-period models and do not allow for the trading of positions in an environment in which liquidity varies from one period to the next. Stretching the period from one day to one week, or to six months, does not make the model more dynamic, as it continues to assume that events occur simultaneously and that the portfolio remains constant during the period. The modeling framework usually does not allow for dynamic hedging or the unwinding of positions.

SUMMARY OF KEY RISKS— VaR AND STRESS TESTING

The stress-testing and scenarios methodologies presented in the previous section can be combined with the VaR approach to produce a summary of significant risks.

Such a report ranks the risk exposure of the firm's positions. For each position, it shows the VaR and the loss corresponding to the stress scenario that would affect the position the most. For example, a high-yield portfolio might well be most exposed to a widening of credit spreads, so the relevant scenario is based on stress-envelope values for a widening of credit spreads.

As stress testing becomes more sophisticated, it's important that institutions devote some thought to how risk analysts report the results of stress tests and present these results to decision makers. There's no point in uncovering the unique vulnerability of an institution to a particular set of events, if that information does not prompt an appropriate response from management.