Value at Risk

Thomas J. Linsmeier and Neil D. Pearson

This article is a self-contained introduction to the concept and methodology of value at risk (VAR), a recently developed tool for measuring an entity's exposure to market risk. We explain the concept of VAR and then describe in detail the three methods for computing it—historical simulation, the delta-normal method, and Monte Carlo simulation. We also discuss the advantages and disadvantages of the three methods for computing VAR. Finally, we briefly describe stress testing and two alternative measures of market risk.

Suppose you are responsible for managing your company's instruments and positions that are sensitive to foreign exchange and interest rate risk. Your boss has been reading about derivative losses suffered by other companies and wants to know if the same thing could happen to his company. That is, he wants to know just how much market risk the company is taking. What do you say?

You could start by listing and describing the company's instruments and positions, but this approach is not likely to be helpful unless the company has only a handful of instruments and positions. Even then, the approach will help only if your boss understands all of the instruments/positions and the risks inherent in each of them.

You could talk about sensitivities (i.e., how much the values of the instruments or positions change when various underlying market rates or prices change and perhaps option deltas and gammas).[1] You are unlikely to win favor with your boss, however, by inundating him with long lists of complex data. In addition, even if you are confident in your ability to explain these data in English, you still have no natural way to net the risk of, for example, your long position in Finnish markka against your short position in Swedish krona.[2]

You could simply assure your boss that you never speculate; you use derivatives only to hedge. He will understand, however, that this statement is vacuous. He knows that the word "hedge" is so ill defined and flexible that virtually any transaction can be characterized as a hedge.

So what do you say?

Thomas J. Linsmeier is associate professor of accounting at Michigan State University. Neil D. Pearson is associate professor of finance at the University of Illinois.

Perhaps the best answer starts: "Our VAR is . . . ."[3] But this answer involves a concept your superiors might never have heard of, let alone understand. This approach does not seem to be a good strategy for getting promoted. The objective of this article, therefore, is to provide a comprehensive introduction to VAR so that you can explain VAR and VAR numbers clearly and simply.

The need for VAR stems from the fact that the past few decades have witnessed tremendous volatility in exchange rates, interest rates, and commodity prices and a proliferation of derivative instruments for use in managing the risks of changes in market rates and prices. The proliferation of derivative instruments has been accompanied by increased trading of securities and a growth of financing opportunities. It has also been coincident with growth in foreign trade and increasing international financial links among companies. As a result of these trends, many companies have portfolios that include large numbers of (sometimes complex) cash and derivative instruments. Because of the complexity and sometimes frequent trading of these instruments, the magnitudes of the risks in companies' portfolios change frequently and often are not obvious. All these trends have led to a demand for a portfolio-level quantitative measure of market risk that a manager can succinctly report to the senior managers charged with the oversight of risk management and trading operations. VAR is the leading summary measure of this type.

The concept and use of VAR is relatively recent. VAR was first used by major financial firms in the late 1980s to measure the risks of their trading portfolios. Since then, the use of VAR has exploded. J.P. Morgan's attempt to establish a market standard through its RiskMetrics™ system in 1994 (J.P. Morgan 1994) provided a tremendous impetus to the growth. VAR is now widely used by other financial institutions, nonfinancial corporations, and institutional investors. Even regulators have become interested in VAR. For example, the Basle Committee on Banking Supervision (1996) permits banks to calculate their capital requirements for market risk using their own proprietary VAR models, and the U.S. Securities and Exchange Commission (1997), which requires that U.S. companies disclose quantitative measures of market risks, lists VAR as one of three possible disclosure methods.

What Is VAR?

VAR is a single, summary statistical measure of possible portfolio losses. VAR is a measure of losses resulting from "normal" market movements. Losses greater than the VAR are suffered only with a specified small probability. Subject to the simplifying assumptions used in its calculation, VAR aggregates all of the risks in a portfolio into a single number suitable for use in the boardroom, reporting to regulators, or disclosure in an annual report. Once one crosses the hurdle of using a statistical measure, the concept of VAR is straightforward to understand. It is simply a way to describe the magnitude of likely losses in a portfolio.

Consider a simple example of VAR involving a foreign exchange (FX) forward contract entered into by a U.S. company at some point in the past. Suppose that the current date is May 20th and the forward contract has 91 days remaining until the delivery date of August 19th. The three-month U.S. dollar and British pound interest rates are, respectively, $r_{USD} = 5.46875$ percent and $r_{GBP} = 6.0625$ percent, and the spot exchange rate is 1.5355 USD/GBP. On the delivery date, the U.S. company will deliver USD15 million and receive GBP10 million. The USD mark-to-market value of the forward contract can be computed by using the interest and exchange rates prevailing on May 20th. Specifically,

USD mark-to-market value

$$= \left[(\text{Exchange rate in USD/GBP}) \times \frac{\text{GBP10 million}}{1 + r_{GBP}(91/360)}\right] - \frac{\text{USD15 million}}{1 + r_{USD}(91/360)}$$

$$= \left[(1.5355\ \text{USD/GBP}) \times \frac{\text{GBP10 million}}{1 + 0.060625(91/360)}\right] - \frac{\text{USD15 million}}{1 + 0.0546875(91/360)}$$

$$= \text{USD327,771}.$$

This calculation uses the fact that one leg of the forward contract is equivalent to a GBP-denominated 91-day zero-coupon bond and the other leg is equivalent to a USD-denominated 91-day zero-coupon bond.

On the next day, May 21st, interest rates, exchange rates, and thus the value of the forward contract are all likely to have changed. Suppose that the distribution of possible one-day changes in the value of the forward contract is the distribution shown in **Figure 1**. The figure indicates that the probability that the loss will exceed USD130,000 is 2 percent, the probability that the loss will be between USD110,000 and USD130,000 is 1 percent, and the probability that the loss will be between USD90,000 and USD110,000 is 2 percent. Summing these probabilities produces a 5 percent probability that the loss will exceed approximately USD90,000. If you deem a loss that is suffered less than 5 percent of the time to be a loss resulting from unusual or "abnormal" market movements, then USD90,000 divides the losses between those from "normal" market movements and those from abnormal movements. Using this 5 percent probability as the cutoff means that USD90,000 is the (approximate) VAR.[4]

The probability used as the cutoff need not be 5 percent; it is chosen by either the user or the provider of the VAR number—perhaps the risk manager, risk-management committee, or designer of the system used to compute the VAR. If the probability chosen were 2 percent, the VAR would be USD130,000 because the loss is predicted to exceed USD130,000 only 2 percent of the time.

Also implicit in this discussion is a choice of holding period. Figure 1 displays the distribution of *daily* profits and losses. One could construct a similar distribution of 5-day or 10-day profits and losses or even use a longer time horizon. Because 5- or 10-day profits and losses typically are larger than 1-day profits and losses, the distributions would be more dispersed or spread out if the longer horizon were chosen and the loss that is exceeded only 5 percent (or 2 percent) of the time would be larger. So, the VAR would be larger.

Now that you have seen an example of VAR, we can proceed to the definition: With a probability of x percent and a holding period of t days, an entity's VAR is the loss that is expected to be exceeded with a probability of only x percent during the next t-day holding period. Alternatively, VAR is the loss that is expected to be exceeded during x percent of t-day holding periods. Typical values for the probability x are 1, 2.5, and 5 percent; common holding periods are 1, 2, and 10 (business) days and 1 month.

Theory provides little guidance about the choice of x. It is determined primarily by how the designer of the risk-management system wants to interpret the VAR number: Is an abnormal loss one that occurs with a probability of 1 percent or 5 percent? For example, the RiskMetrics system uses

Figure 1. Histogram of Hypothetical Daily Mark-to-Market Profits and Losses on a Forward Contract

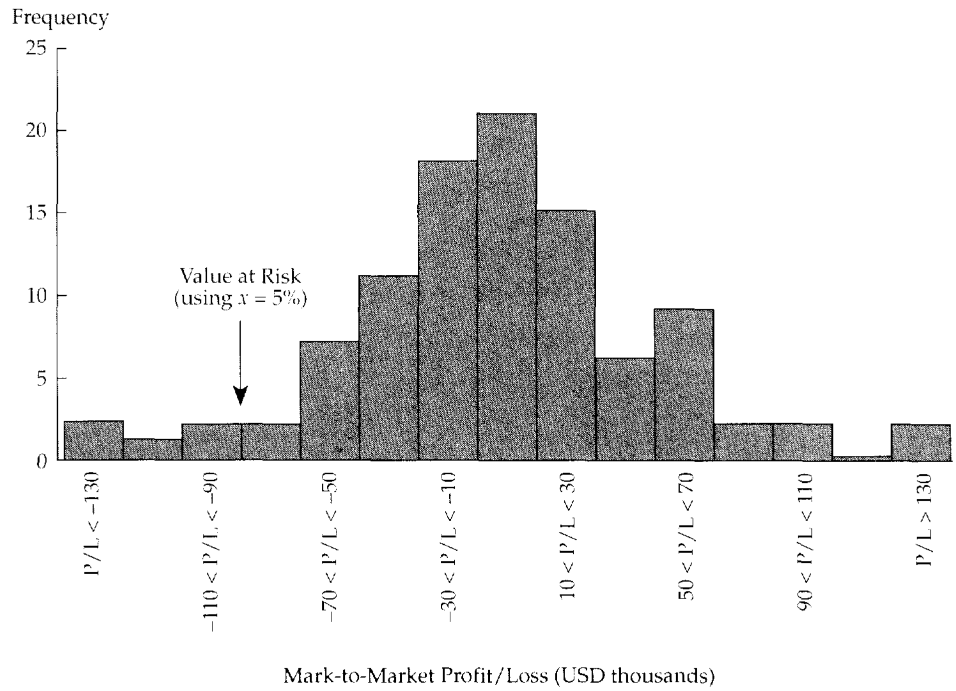

5 percent whereas Mobil Oil's 1998 annual report indicates that Mobil uses 0.3 percent.

The parameter t is determined by the entity's horizon. Those that actively trade their portfolios, such as financial firms, typically use one day, whereas institutional investors and nonfinancial corporations may use longer periods. A VAR number applies to the current portfolio, so one (sometimes implicit) assumption underlying the computation is that the current portfolio will remain unchanged throughout the holding period. This assumption may not be reasonable for long holding periods.

In interpreting VAR numbers, one must keep in mind the probability x and holding period t. Without them, VAR numbers are meaningless. Two companies holding identical portfolios will come up with different VAR estimates if they make different choices of x and t. Obviously, the loss that is suffered with a probability of only 1 percent will be a larger amount than the loss that is suffered with a probability of 5 percent. Under the assumptions used in some VAR systems, it is 1.414 times as large.[5] The choice of holding period can have an even larger impact because the VAR computed using a t-day holding period is approximately $\sqrt{t}$ times as large as the VAR computed using a one-day holding period. In the absence of appropriate adjustments for these factors, VAR numbers are not comparable across entities.

Identifying the Important Market Factors

To compute VAR (or any other quantitative measure of market risk), one needs to identify the basic market rates and prices that affect the value of the portfolio—that is, the market factors. Identifying a limited number of basic market factors is necessary because, otherwise, the problem of computing a portfolio-level quantitative measure of market risk becomes unmanageable. Even for simple instruments, such as forward contracts, an almost countless number of contracts are possible because virtually any forward price and delivery date are possible. Other instruments, such as swaps, loans with embedded options, options, and exotic options, are even more complicated. Thus, expressing the instruments' values in terms of a limited number of basic market factors is an essential first step.

Typically, market factors are identified by decomposing the instruments in the portfolio into simpler instruments more directly related to basic market risk factors and then interpreting the actual instruments as portfolios of the simpler instruments. We will illustrate this concept using the FX forward contract we introduced earlier.

So, the current date is May 20th. The contract requires a U.S. company to deliver USD15 million in 91 days. In exchange, it will receive GBP10 million. The current USD market value of this forward contract depends on three basic market factors: S, the spot exchange rate expressed in U.S. dollars per British pound; r_{GBP}, the three-month British pound interest rate; and r_{USD}, the three-month U.S. dollar interest rate. To explain, we decompose the cash flows of the forward contract into a portfolio consisting of a long position in a 91-day zero-coupon bond with a face value of GBP10 million and a short position in a 91-day zero-coupon bond with a face value of USD15 million. This decomposition yields the following formula, which was used in the opening example, for the current mark-to-market value (in U.S. dollars) of the position in terms of the basic market factors, r_{USD}, r_{GBP}, and S:

$$\text{USD mark-to-market value} = \left[S \times \frac{\text{GBP10 million}}{1 + r_{GBP}(91/360)}\right] - \frac{\text{USD15 million}}{1 + r_{USD}(91/360)}. \quad (1)$$

Similar formulas expressing the instruments' values in terms of the basic market factors must be obtained for all of the instruments in the portfolio.[6] Once the formulas have been obtained, a key part of the problem of quantifying market risk is finished. The remaining steps involve determining or estimating the statistical distribution of the potential future values of the market factors, using these future market factors and formulas to determine potential future changes in the values of the various positions that compose the portfolio, and then aggregating across positions in order to determine the potential future changes in portfolio value. VAR is a measure of these potential changes in portfolio value.

Of course, the values of most actual portfolios depend on more than three market factors. A typical set of market factors might include the spot exchange rates for all currencies in which the company has positions, together with (for each currency) the interest rates on zero-coupon bonds for a range of maturities. For example, the maturities used in the first version of the RiskMetrics system were 1 day, 1 week, 1, 3, 6, and 12 months, and 2, 3, 4, 5, 7, 9, 10, 15, 20, and 30 years.[7] A company with positions in most interest rates and actively traded currencies can easily have a portfolio exposed to several hundred market factors.

The dependence of VAR on only a limited number of basic market factors often remains implicit in the historical and Monte Carlo simulation methodologies, but it must be made explicit in the delta-normal methodology. The process of making the dependence explicit is known as "risk mapping," and we describe it in detail in a later section.

VAR Methodologies

The three basic methods of calculating VAR are historical simulation, the delta-normal approach,[8] and Monte Carlo simulation.

Historical Simulation. Historical simulation requires relatively few assumptions about the statistical distributions of the underlying market factors. We illustrate the procedure with a simple portfolio consisting of the single three-month FX forward contract described previously, and we perform the analysis from the perspective of the U.S. company. Even though our example is of a single-instrument portfolio, it captures some of the features of multiple-instrument portfolios because the forward contract is exposed to the risk of changes in several basic market factors.

In essence, the approach involves using historical changes in market rates and prices to construct a distribution of potential future portfolio profits and losses (similar to Figure 1) and then reading off the VAR as the loss that is exceeded only 5 percent of the time. The distribution of profits and losses is constructed by taking the *current* portfolio and subjecting it to the *actual* changes in the market factors experienced during each of the last N periods (periods here being days). That is, N sets of hypothetical market factors are constructed from their current values and the changes experienced during the last N periods. Using these hypothetical values of the market factors, N hypothetical mark-to-market portfolio values are computed. This step allows one to compute N hypothetical mark-to-market profits and losses on the portfolio, as compared with the current mark-to-market portfolio value. (Even though we use the actual changes in rates and prices, the mark-to-market profits and losses are hypothetical because the current portfolio was not held in each of the last N periods.) The use of actual historical changes in rates and prices to compute the hypothetical profits and losses is the distinguishing feature of historical simulation. Once the hypothetical mark-to-market profit or loss for each of the last N periods have been calculated, the distribution of profits and losses and the VAR can be determined.

Single-instrument portfolio. Recall that the current date is May 20th and that the forward contract obligates a U.S. company to deliver USD15 million on the delivery date 91 days hence and receive in exchange GBP10 million. The holding period is one day ($t = 1$), the VAR will be computed using a 5 percent probability ($x = 5$ percent), and the most recent 100 business days ($N = 100$) will be used to compute the changes in the values of the market factors and the hypothetical profits and losses on the portfolio. Because in our example May 20th is the 100th business day of the fiscal year, the most recent 100 business days start on January 2.

Historical simulation is carried out in five steps.

Step 1. The first step is to identify the basic market factors and obtain a formula expressing the mark-to-market value of the forward contract in terms of the market factors. The market factors were identified in the previous section: the three-month British pound interest rate, the three-month U.S. dollar interest rate, and the spot USD/GBP exchange rate. Also, we previously derived a formula for the dollar mark-to-market value of the forward contract by decomposing it into a long position in a pound-denominated zero-coupon bond with face value of GBP10 million and short position in a dollar-denominated zero-coupon bond with face value of USD15 million.

Step 2. The next step is to obtain historical values of the market factors for the last N periods. For the example portfolio, we need to collect the three-month U.S. dollar and British pound interest rates and the spot USD/GBP exchange rate for the last 100 business days. Daily changes in these rates will be used to construct hypothetical values of the market factors for the calculation of hypothetical profits and losses in Step 3, because the daily VAR number measures the hypothetical portfolio loss over a one-day holding period (e.g., May 20th to May 21st).

Step 3. This is the key step. We subject the portfolio to the changes in market rates and prices experienced on each of the most recent 100 business days by calculating the daily profits and losses that would have occurred if comparable daily changes in the market factors had been experienced and the *current* portfolio had been marked to market.

To calculate the 100 daily profits and losses, we first calculate 100 sets of hypothetical values of the market factors. The hypothetical market factors are based on, but *not* equal to, the historical values of the market factors over the past 100 days. Instead of using actual historical values, we calculate daily historical percentage changes in the market factors and then combine the historical percentage changes with the current (May 20th) market factors to compute 100 sets of hypothetical market factors on May 21st.[9] We will use these hypothetical market factors to calculate the 100 hypothetical mark-to-market portfolio values. For each of the hypothetical portfolio values, we subtract the actual mark-to-market portfolio value on May 20th to obtain 100 hypothetical daily profits and losses.

Table 1 shows the calculation of the hypothetical profits and losses based on the changes in the market factors from the first business day of the current year, which is Day 1 of the 100 days preceding May 20th. We start by using the May 20th values of the market factors to compute the mark-to-market value of the forward contract on May 20th, which is shown on Line 1. Next, we determine what the value might be on the next day. To do this, we use the percentage changes in the market factors from December 29th of the prior year to January 2nd of the current year. The actual values on December 29th and January 2nd and the percentage changes are shown on Lines 2 through 4. Then, on Lines 5 and 6, we show the result of using the values of the market factors on May 20th with the percentage changes from December 29th to January 2nd to compute hypothetical values of the market factors for

Table 1. Calculation of Hypothetical May 21st Mark-to-Market Profit/Loss on Forward Contract: Historical Simulation Method

	Market Factors			
Measure of Value	USD Interest Rate (% per year)	GBP Interest Rate (% per year)	Exchange Rate (USD/GBP)	Mark-to-Market Value of Forward Contract (USD)
1. Actual values as of close of business on 5/20	5.46875%	6.0625%	1.5355	$327,771
2. Actual values on 12/29/prior year	5.6875	6.5000	1.5530	
3. Actual values on 1/2/current year	5.6875	6.5625	1.5568	
4. Percentage change 12/29 to 1/2	0.0000	0.9620	0.2430	
5. Actual values on 5/20	5.46875	6.0625	1.5355	327,771
6. Hypothetical future values calculated using rates from 5/20 and percentage changes from 12/29 to 1/2	5.46875	6.1208	1.5392	362,713
7. Hypothetical mark-to-market profit/loss on forward contract				34,942

Note: The hypothetical future value of the forward contract is computed using the formula

$$\text{USD mark-to-market value} = \text{Exchange rate in USD/GBP} \times \frac{\text{GBP10 million}}{1 + r_{GBP}(90/360)} - \frac{\text{USD15 million}}{1 + r_{USD}(90/360)}\ \text{USD}.$$

May 21st. We then use these hypothetical values on May 21st and Equation 1 to compute a mark-to-market value of the forward contract for May 21st. This value is also shown on Line 6. Once the hypothetical May 21st mark-to-market value has been computed, the profit or loss on the forward contract is simply the change in the mark-to-market value from May 20th to May 21st, shown on Line 7.

We use the values of the market factors on May 20th and the percentage changes in the market factors for Days 2 through 100 to repeat this calculation 99 more times. In doing so, we compute 100 hypothetical mark-to-market values of the forward contract for May 21st and 100 hypothetical mark-to-market profits or losses.

Step 4. The next step is to order the mark-to-market profits and losses from the largest profit to the largest loss. The ordered profits/losses, the beginning and end of which are shown in **Table 2**, range from a profit of USD212,050 to a loss of USD143,207.

Step 5. Finally, we select the loss that is equaled or exceeded 5 percent of the time. Because we are using 100 days, this loss is the fifth worst loss, or the loss of USD97,230 (in the box in Table 2). Using a probability of 5 percent, this number is the VAR. In Figure 1, the VAR, the loss that leaves 5 percent of the probability in the left-hand tail, is indicated by an arrow.

Multi-instrument portfolios. Using a three-month forward contract in the example allowed us to sidestep one minor difficulty that arises when multiple-instrument portfolios are being considered: If the market risk factors include the spot exchange rates and the interest rates at 1, 3, 6, and 12 months, what do we do with a 4-month forward contract? The obvious answer is to write a formula for the contract value in terms of the four-month U.S. dollar and British pound interest rates and the USD/GBP exchange rate, just as we did with the three-month forward. But doesn't this approach introduce two more market factors—the four-month dollar and pound interest rates?

The answer is no—but multi-instrument portfolios do require an extension of the method. The 1-, 3-, 6-, and 12-month interest rates are natural choices for market risk factors because interbank

Table 2. 100 Hypothetical Daily Mark-to-Market Profits and Losses Ordered from Largest Profit to Largest Loss: Historical Simulation

	Market Factors			Hypothetical Mark-to-Market Value of Forward Contract (USD)	Change in Mark-to-Market Value of Forward Contract (USD)
Number	USD Interest Rate (% per year)	GBP Interest Rate (% per year)	Exchange Rate (USD/GBP)		
1	5.469%	6.063%	1.557	$539,821	$212,050
2	5.469	6.063	1.551	480,897	153,126
3	5.469	6.063	1.546	434,228	106,457
4	5.469	6.063	1.545	425,982	98,211
5	5.532	6.063	1.544	413,263	85,492
6	5.532	6.126	1.543	398,996	71,225
7	5.469	6.063	1.542	396,685	68,914
8	5.469	6.063	1.542	392,978	65,207
9	5.469	6.063	1.542	392,571	64,800
10	5.469	6.063	1.541	385,563	57,792
•					
•					
•					
91	5.469	6.005	1.529	270,141	–57,630
92	5.500	6.063	1.529	269,264	–58,507
93	5.531	6.063	1.529	267,692	–60,079
94	5.469	6.004	1.528	255,632	–72,139
95	5.469	6.063	1.528	249,295	–78,476
96	5.469	6.063	1.526	230,541	–97,230
97	5.438	6.063	1.526	230,319	–97,452
98	5.438	6.063	1.523	203,798	–123,973
99	5.438	6.063	1.522	196,208	–131,563
100	5.407	6.063	1.521	184,564	–143,207

deposit markets are active at these maturities and rates for these maturities are widely quoted. In addition, for a number of currencies, there are also liquid government bond markets at some of these maturities. But typically, four-month interbank markets are not active. As a result, the four-month interest rates used in computing the model value of the four-month forward contract are typically interpolated from the three-month and six-month interest rates. (The interpolated four-month rates might also depend on rates for the other actively quoted maturities, depending on the interpolation scheme used.) Thus, the current mark-to-market values of all USD/GBP forward contracts, regardless of delivery date, will depend on the spot exchange rate and the interest rates at only a limited number of maturities.

Extending the methodology to handle realistic multi-instrument portfolios also requires a bit of additional work in three of the steps. First, Step 1 is likely to involve many more market factors, namely, the interest rates for longer maturity bonds and the interest and exchange rates for other currencies. These factors must be identified, and pricing formulas expressing the instruments' values in terms of the market factors must be obtained. Options may be handled either by treating the option volatilities as additional market factors that must be estimated and collected in each of the last N periods or by treating the volatilities as constants and disregarding that they change randomly over time. Second, in Step 2, the historical values of all the market factors must be collected. Third, a crucial aspect is that the mark-to-market profits and losses on each instrument in the portfolio be computed and then summed for each day before they are ordered from highest profit to lowest loss in Step 4. The calculation of VAR is intended to capture the fact that gains on some instruments typically offset losses on others. Netting the gains against the losses for each of the 100 days in Step 3 captures these correlations.

Delta-Normal Approach. The delta-normal approach is based on the assumption that the underlying market factors have a multivariate normal distribution. Using this assumption (and the approximation involved in mapping the portfolio that is detailed later), one can determine the distribution of mark-to-market portfolio profits and losses, which is also assumed to be normal. Once the distribution of possible portfolio profits and losses has been obtained, standard mathematical properties of the normal distribution are used to determine the loss that will be equaled or exceeded x percent of the time (i.e., the VAR).

To illustrate, we continue with the three-month FX forward contract introduced earlier and continue to assume that the holding period is one day.

For the normal distribution, outcomes less than or equal to 1.65 standard deviations below the mean occur only 5 percent of the time. Because the VAR is defined as the loss that is exceeded with a probability of x percent, the VAR is

$$\text{Value at risk} = -\left[\begin{pmatrix}\text{Expected change}\\ \text{in portfolio value}\end{pmatrix} - 1.65\begin{pmatrix}\text{Standard deviation of}\\ \text{change in portfolio value}\end{pmatrix}\right]. \quad (2)$$

For this example, we assume that the expected change in portfolio value is zero. This assumption is reasonable for a short holding period and is frequently made. The standard deviation, which is a measure of the "spread" or dispersion of the distribution, is approximately USD52,500. This normal distribution with a mean of zero and standard deviation of 52,500 can be represented by the probability density function shown in **Figure 2**. Using this distribution, we find the VAR is –[0 – 1.65(USD52,500)] = USD86,625.

Clearly, computation of the standard deviation of changes in portfolio value is the key step in the delta-normal approach. The standard deviation of the changes in value of a portfolio depends on the standard deviations of the changes in value of all the instruments in the portfolio and the correlations among them; with n instruments in the portfolio, the total is $n(n + 1)/2$ parameters. But estimating all of these standard deviations and correlations directly is usually not feasible. The necessary data simply may not exist, either because the securities or instruments have only recently been traded or because they have only recently been added to the portfolio. Even when some data are available, reliable estimation of all the standard deviations and correlations requires that the number of observations be considerably greater than the number of instruments n, and so much data will be available only when n is relatively small.

As a result, most risk-measurement systems replace the actual portfolio with a simpler portfolio that has approximately the same risk and then compute the VAR of the simpler portfolio. In particular, they identify a set of k market factors (e.g., yields on zero-coupon bonds) that account for most of the changes in value of the portfolio; then, for each market factor, they identify a standardized position (e.g., a zero-coupon bond) that is exposed to the risk of only one market factor. This "approximating" portfolio has the same exposures to the basic market factors as the original portfolio and, therefore, captures its risks.

Figure 2. Probability Density Function and VAR Obtained from Delta-Normal Method

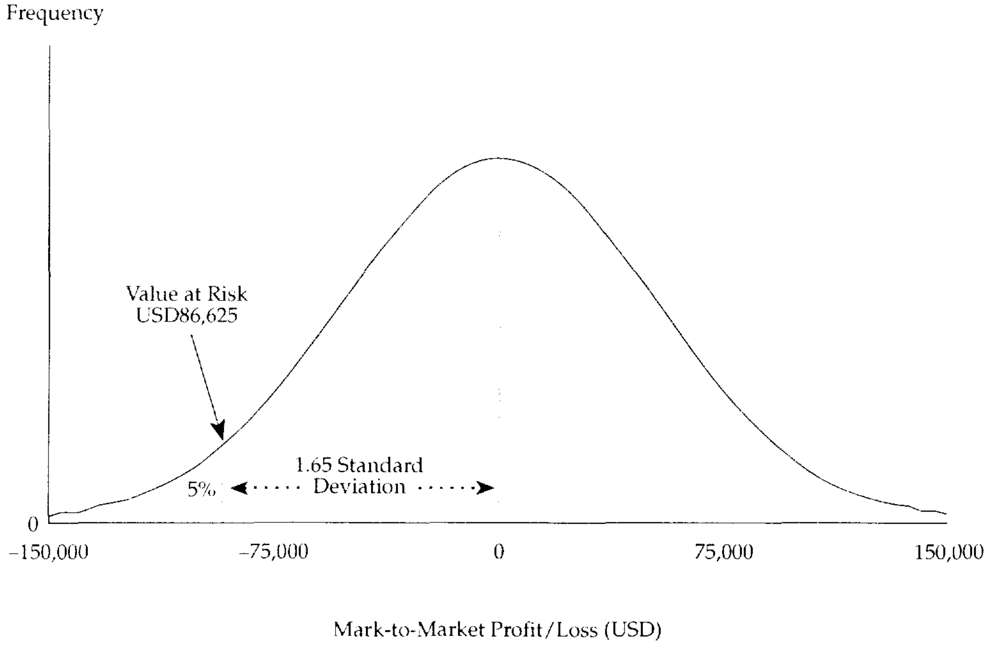

The representation of the portfolio in terms of k standardized positions is known as risk mapping. It is specifically designed to minimize the burden of computing the portfolio standard deviation by limiting the number of standard deviations and correlations in the computation to $k(k+1)/2$. Computational and data costs are minimized in risk mapping because by carefully selecting the market factors, the designer of the risk-measurement system can often make k considerably less than n. In addition, in this approach, problems of data collection are minimized because standard deviations and correlations are computed from time series of changes in market factors, which are more readily obtainable than changes in the values of specific instruments and positions.

In the next two sections and Appendixes A and B, we discuss key issues in the risk-mapping process and provide specific steps for computing the portfolio standard deviation and VAR in the delta-normal approach.

Single-instrument portfolio. The delta-normal method requires four steps to analyze a single-instrument portfolio.

Step 1. The first step is to identify the basic market factors and the standardized positions directly related to the market factors and map the forward contract onto the standardized positions.

The designer of the risk-measurement system has considerable flexibility in the choice of basic market factors and standardized positions and, therefore, considerable flexibility in setting up the risk mapping. The choice of market factors is, in fact, the key issue in designing a risk-measurement system because the choice of market factors amounts to the choice of what risks to measure.[10]

We use a simple set of standardized positions to illustrate the procedure. The natural choice corresponds to our previous decomposition of the forward contract into a long position in a three-month British-pound-denominated zero-coupon bond with a face value of GBP10 million and short position in a three-month U.S.-dollar-denominated zero-coupon bond with a face value of USD15 million. Thus, we take the standardized positions to be three-month dollar-denominated zero-coupon bonds, three-month pound-denominated zero- coupon bonds exposed only to changes in the pound interest rate (i.e., as if the exchange rate were fixed), and a spot position in pounds.

By decomposing the forward contract into a U.S. dollar leg and a British pound leg, we have already completed a good bit of the work involved in mapping the contract. We need only finish the process by specifying the dollar amount of the standardized positions exposed to each market factor.

The U.S. dollar leg of the forward contract is exposed to changes in U.S. interest rates. This leg is equivalent to a short position in a dollar-denominated zero-coupon bond with a face value of USD15 million, and the dollar magnitude of exposure to U.S. interest rates can be obtained by discounting the face value of the contract using the three-month dollar interest rate.

Letting X_1 denote the first standardized position, we compute the exposure to U.S. dollar interest rates to be

$$X_1 = \frac{\text{USD15 million}}{1 + r_{USD}(91/360)}$$
$$= \frac{\text{USD15 million}}{1 + 0.05469(91/360)}$$
$$= -\text{USD14,795,000}.$$

We use a negative sign to represent the short position.

The British pound leg must be mapped into two standardized positions, X_2 and X_3, because its value depends on two market factors—the three-month pound interest rate and the spot USD/GBP exchange rate. The magnitudes of the standardized positions are determined separately by considering how changes in each of the market factors affect the U.S. dollar value of the pound leg while holding the other factor constant. We have

$$\text{USD value of GBP leg} = \frac{\text{GBP10 million}}{1 + r_{GBP}(91/360)} \times (\text{USD } S/\text{GBP})$$
$$= \frac{\text{GBP10 million}}{1 + 0.06063(91/360)} \times (\text{USD } 1.5355/\text{GBP})$$
$$= \text{USD15,123,242}.$$

Holding the spot exchange rate S constant, the risk of X_2 is USD15,123,242 invested in three-month British pound bonds. Holding the pound interest rate constant, the bond with a face value of GBP10 million has the exchange rate risk of a spot position (its present value) of [GBP10 million/[1 + 0.06063(91/360)] pounds or USD15,123,242. Hence, the dollar value of the spot pound position is X_3 = USD15,123,242.

The equality of X_2 and X_3 is not a coincidence; both represent the U.S. dollar value of the British pound leg of the forward contract. This value appears twice in the mapped position because, from the perspective of a U.S. company, a position in a pound-denominated bond is exposed to changes in two market risk factors.

When this mapping is complete, the forward contract is now described by the magnitudes of three standardized positions: X_1 = USD – 14,795,000, X_2 and X_3 = USD15,123,242. Appendix A provides a mathematical argument that justifies this mapping.

Step 2. The second step is to assume that percentage changes in the basic market factors have a multivariate normal distribution and then to estimate the parameters of that distribution. The estimated standard deviations and correlation coefficients are in **Table 3**. For this example, we assume that the means are zero.

Step 3. The next step is to use the standard deviations and correlations of the market factors to determine the standard deviations and correlations of changes in the values of the standardized positions. The standard deviations of changes in the values of the standardized positions are determined by the products of the standard deviations of the market factors and the sensitivities of the standardized positions to changes in the market factors. For example, if the value of the first standardized position changes by 2 percent when the first market factor changes by 1 percent, then its standard deviation is twice as large as the standard deviation of the market factor. The correlations between changes in the values of standardized positions are equal to the correlations between the market factors, except that the correlation coefficient changes sign if the value of one of the standardized positions changes inversely with changes in the corresponding market factor. For example, Table 3 indicates that the correlation between the U.S. dollar interest rate (the first market factor) and the USD/GBP exchange rate (the third market factor) is 0.19. Thus, the correlation between the values of the first and third standardized positions is –0.19 because the value of the first standardized position moves inversely with changes in the third standardized position.

Step 4. Now that we have the standard deviations of and correlations between changes in the values of the standardized positions, we can use standard mathematical results about the distributions of sums of normal random variables to calculate the portfolio variance and standard deviation,

Table 3. Standard Deviations of and Correlations between Percentage Changes in Market Factors

Market Factor	Standard Deviation of Percentage Changes	Correlations between Percentage Changes in Market Factors: Three-Month USD Interest Rate	Three-Month GBP Interest Rate	USD/GBP Exchange Rate
Three-month USD interest rate	0.61	1.00		
Three-month GBP interest rate	0.58	0.11	1.00	
USD/GBP exchange rate	0.35	0.19	0.10	1.00

and we can determine the distribution of portfolio profit or loss. The variance of changes in mark-to-market portfolio value depends on the standard deviations, σ, of changes in the value of the standardized positions;[11] the correlations, ρ; and the sizes of the positions, X. The variance of changes in mark-to-market portfolio value is given by the formula

$$\sigma^2_{portfolio} = X_1^2\sigma_1^2 + X_2^2\sigma_2^2 + X_3^2\sigma_3^2 + 2X_1X_2\rho_{1,2}\sigma_1\sigma_2 + 2X_1X_3\rho_{1,3}\sigma_1\sigma_3 + 2X_2X_3\rho_{2,3}\sigma_2\sigma_3. \quad (3)$$

For our example, the portfolio standard deviation is approximately $\sigma_{portfolio}$ = USD52,5000. Using Equation 2 and the assumption that the expected change in the portfolio value is zero, we can compute the value at risk as

$$\begin{aligned}\text{VAR} &= [0 - (1.65 \times \sigma_{portfolio})] \\ &= \text{USD86,625}.\end{aligned}$$

Figure 2 shows the probability density function for a normal distribution with a mean of zero and a standard deviation of USD52,500, together with the VAR.

Multi-instrument portfolios. Multi-instrument portfolios are handled by mapping each of the instruments to the standardized positions and then computing the VAR of the aggregate portfolio of standardized positions. Of course, a multi-instrument portfolio will have many more market factors and standardized positions than a single instrument. More importantly, the mapping is more complicated.

We can examine the multi-instrument challenge by starting with the question discussed for the historical simulation method: How would a four-month instrument be handled in the delta-normal method? In this approach, a four-month forward would be handled by first decomposing the contract into British-pound-denominated and U.S.-dollar-denominated four-month zero-coupon bonds (just as in the three-month forward). Next, the four-month zeros would be mapped onto the three-month and six-month zeros. The idea is to replace each of the four-month zero-coupon bonds with a portfolio of three-month and six-month standardized positions with the same market value and risk (in this context, "risk" might mean either standard deviation of changes in mark-to-market value or duration). An instrument with multiple cash flows at different dates—for example, a 10-year British government bond—would be handled by mapping the 20 semiannual cash flows onto the 3-month and 6-month and 1-, 2-, 3-, 4-, 5-, 7-, and 10-year pound-denominated zero-coupon bonds, which are the standardized positions. Each cash flow would be mapped onto the two nearest standardized positions.

Forward contracts and bonds decompose naturally into portfolios of zero-coupon bonds and thus map naturally into the standardized positions. This is not the case for options and other securities or instruments with embedded options. To map such instruments, we need to rely on a more general approach, one that applies to all instruments.

This general approach is based precisely on the theoretical framework discussed in Appendix A. In this approach, to measure the risk of an option or other instrument, one uses its sensitivities, or deltas (partial derivatives), with respect to each of the market factors. The deltas indicate the magnitudes of the changes in the value of the instrument resulting from changes in each of the market factors. Thus, the deltas can be used to measure the risk of an option or another instrument; two instruments or portfolios with the same deltas have the same risk. Using this idea, options and other instruments are mapped to a portfolio of standardized positions with the same deltas or exposures to the basic market factors. Appendix B discusses this procedure in more detail. The mapping of the forward contract is actually no more than a special case of this approach.

An important limitation of the delta-normal approach is that it maps instruments to their delta-equivalent positions (i.e., it amounts to finding a portfolio of standardized positions that have the same deltas as the original instruments), but if the instruments' deltas change as the market factors change (as is common with options), delta will not completely measure the changes in value. Appendix B also discusses this issue.

Monte Carlo Simulation. The Monte Carlo simulation methodology has a number of similarities to historical simulation. The main difference is that, rather than carrying out the simulation using the observed changes in the market factors over the last N periods to generate N hypothetical portfolio profits or losses, in Monte Carlo simulation, one chooses a statistical distribution that is believed to adequately capture or approximate the possible changes in the market factors. Then, a pseudorandom number generator is used to generate thousands (or perhaps tens of thousands) of hypothetical changes in the market factors. These hypothetical changes are used to construct thousands of hypothetical portfolio profits and losses on the current portfolio and the distribution of possible portfolio profit or loss. Finally, the VAR is determined from this distribution.

Single-instrument portfolio. We use the same forward contract to illustrate the approach. The steps are as follows:

Step 1. The first step is to identify the basic market factors and obtain a formula expressing the mark-to-market value of the forward contract in terms of the market factors. We did this step in the discussion of historical simulation: The market factors are the three-month British pound interest rate, the three-month U.S. dollar interest rate, and the spot USD/GBP exchange rate. We have already derived a formula (Equation 1) for the mark-to-market value of the forward contract by decomposing it into a portfolio of dollar- and pound-denominated three-month zero-coupon bonds.

Step 2. The second step is to determine or assume a specific distribution for changes in the basic market factors and estimate the parameters of that distribution. The ability to pick the distribution is the feature that distinguishes Monte Carlo simulation from the other two approaches, for in the other two methods, the distribution of changes in the market factors is specified as part of the method. For this example, we assume that percentage changes in the basic market factors have a multivariate normal distribution and use the estimates of the standard deviations and correlations in Table 3.

The assumed distribution need not be multivariate normal.[12] The designers of the risk-management system are free to choose any distribution they think reasonably describes possible future changes in the market factors. Beliefs about future changes in the market factors are typically based on observed past changes, so in effect, the designers of the risk-management system are free to chose any distribution they think approximates the distribution of past changes.

Step 3. Once the distribution has been selected, the next step is to use a pseudo-random generator to generate N hypothetical values of changes in the market factors, where N is almost certainly greater than 1,000 and perhaps greater than 10,000. These hypothetical market factors are then used to calculate N hypothetical mark-to-market portfolio values. Finally, from each of the hypothetical portfolio values, one subtracts the actual mark-to-market portfolio value on May 20th to obtain N hypothetical daily profits and losses.

Steps 4 and 5. The last two steps are the same as in historical simulation. The mark-to-market profits and losses are ordered from the largest profit to the largest loss, and the VAR is the loss that is equaled or exceeded 5 percent of the time.

Multi-instrument portfolios. As with historical simulation, extending the methodology to handle realistic multi-instrument portfolios requires only that some additional work be performed in three of the steps. First, many more market factors are likely than in a single-instrument analysis—in particular, the interest rates for longer maturity bonds and the interest and exchange rates for other currencies. These factors must be identified, and pricing formulas expressing the instruments' values in terms of the market factors must be obtained. Again, as in historical simulation, options may be handled either by treating the option volatilities as additional market factors that must be simulated or treating the volatilities as constants and disregarding that they change randomly over time. Second, the joint distribution of possible changes in the values of all the market factors must be determined. This joint distribution must include the option volatilities if they are to be allowed to change. Third, similar to historical simulation, to reflect accurately the correlations of market rates and prices, the mark-to-market profits and losses on every instrument must be computed and then summed for each day before they are ordered from highest profit to lowest loss.

Which Method Is Best?

Unfortunately, no easy answer exists to the question of which method of calculating VAR is best. The methods differ in ability to capture the risks of options and option-like instruments, ease of implementation, ease of explanation to senior managers, flexibility in analyzing the effect of changes in the assumptions, and reliability of results. The best choice will be determined by which dimensions the risk manager considers most important. In the next sections, we discuss how the three methods differ along these dimensions, and **Exhibit 1** summarizes the differences. We also discuss a closely related issue, namely, the choice of holding period.

Capturing the Risks of Options and Option-Like Instruments. The two simulation methods work well regardless of the presence of options and option-like instruments in the portfolio. In contrast, the standard delta-normal method works well for instruments and portfolios with little option content but not as well as the two simulation methods when options and option-like instruments are significant in the portfolio. The limitation of the delta-normal method is that it typically incorporates options by replacing them with or mapping them to their "delta-equivalent" spot positions. This process amounts to linearizing the option positions or replacing the nonlinear functions that give their values with linear approximations. For instruments or portfolios with a great deal of option

Exhibit 1. Comparison of VAR Methodologies

Attribute	Historical Simulation	Delta-Normal	Monte Carlo Simulation
Is it able to capture the risks of portfolios that include options?	Yes, regardless of the option content of the portfolio	No, except when computed using a short holding period for portfolios with limited or moderate option content	Yes, regardless of the option content of the portfolio
Is it easy to implement?	Yes for portfolios for which data on the past values of the market factors are available	Yes for portfolios restricted to instruments and currencies covered by available off-the-shelf software; otherwise, reasonably easy to moderately difficult to implement, depending on the complexity of the instruments and availability of data	Yes for portfolios restricted to instruments and currencies covered by available off-the-shelf software; otherwise, moderately to extremely difficult to implement
Are the computations performed quickly?	Yes	Yes	No, except for relatively small portfolios
Is it easy to explain to senior managers?	Yes	No	No
Does it produce misleading VAR estimates when recent past is atypical?	Yes	Yes, except that alternative correlations/standard deviations may be used	Yes, except that alternative estimates of parameters may be used
Are "what-if" analyses to examine effect of alternative assumptions easy to perform?	No	Examining alternative assumptions about correlations/standard deviations is easy; examining alternative assumptions about the distribution of the market factors (i.e. distributions other than the normal) is impossible	Yes

content, the linear approximations may not adequately capture how the values of the options change with changes in the underlying rates and prices (see Appendix B).

In the delta-normal method, the problem of adequately capturing the risks of options and option-like instruments is least severe when the holding period is one day. Large changes in the underlying rates or prices are unlikely over such a short period, and the linear approximation in this method works well for small changes in rates and prices. Over longer holding periods—one month, for example—larger changes in underlying rates and prices are likely and VAR estimates produced by the delta-normal method cannot be relied on for positions with moderate or significant option content.

The simulation methods work well when options are in the portfolio because they recompute the value of the portfolio for each "draw" of the basic market factors. In doing so, they estimate the "correct" distribution of portfolio value. This statement must be qualified, however, because the distribution of portfolio value generated by Monte Carlo simulation depends on the assumed statistical distribution of the basic market factors and estimates of its parameters. To the extent that a poor choice of distribution is made or the estimates of the parameters are poor, the method will lead to errors in the calculated VAR. Similarly, the distribution of portfolio value generated by historical simulation will be misleading if the prior N days from which the historical sample was drawn were not representative.

A final risk-measurement issue related to options and option-like instruments is the ability of the VAR methodologies to incorporate the fact that option volatilities change over time. In principle, the delta-normal and Monte Carlo methods can incorporate changing option volatilities by including them as additional market factors, although this step typically is not taken in practice. Historical simulation also can incorporate changes in option prices with changes in volatilities if option volatilities are included as additional factors and collected for the N-day period used in the simulation.

Ease of Implementation. Historical simulation is easy to implement for portfolios restricted to instruments for which data on the past values of the basic market factors are available. It is conceptually simple, and it can be implemented in a spreadsheet because pricing models for most financial products are available as spreadsheet add-in functions. The principal difficulty in implementing historical simulation is that it requires the user to possess a time series of the relevant market factors covering the last N days, which can pose a problem for multinational companies with operations in many countries or with receivables and other instruments that are sensitive to changes in a wide range of market factors. Although spot exchange rates are readily available for virtually all currencies, obtaining reliable daily market interest rates for a range of maturities in countries without well-developed capital markets can be difficult.

Several vendors offer software that uses the delta-normal method to compute VAR estimates, so this method is easy to implement for portfolios with market factors and types of instruments covered by the available systems. But the delta-normal method can be moderately difficult to implement for portfolios that include market factors and types of instruments not covered by the available systems. First, the method requires estimates of the standard deviations and correlations of the market factors, and although computing these estimates is straightforward (if data are available), reliable market interest rates may not be available for all the maturities in all the currencies represented in the portfolio. A second, and more difficult, problem is that the instruments must be mapped to delta-equivalent positions, as described in Appendixes A and B.

Off-the-shelf software is also available for Monte Carlo simulation, which makes it as easy to implement as the delta-normal method for portfolios covered by the available systems. For portfolios not covered by existing software, Monte Carlo simulation is in some ways easier and in some ways more difficult than the delta-normal method. It is easier because it does not require mapping of the instruments onto standardized positions. It is more difficult because the user must select the distribution from which the pseudo-random vectors are drawn and select or estimate the parameters of that distribution. Actually carrying out the simulation is not difficult because pseudo-random number generators are available as spreadsheet add-ins. Selecting the distribution and selecting or estimating the parameters, however, require great expertise and judgment. Another disadvantage of Monte Carlo simulation is that the computations for large portfolios can be time-consuming.

All three methods require pricing models for all instruments in the portfolio.[13] (Although the delta-normal method does not directly make use of instruments' prices, options are mapped to their delta-equivalent positions and the computation of deltas requires pricing models.) The need for pricing models can pose a problem for complex instruments and portfolios that include such instruments.

Ease of Communication. The conceptual simplicity of historical simulation makes it easiest to explain to senior managers. The delta-normal method is difficult to explain to those who lack technical training because the key step, the mathematics of the normal distribution in calculating portfolio standard deviation and VAR, is simply a black box. Monte Carlo simulation is even more difficult to explain. Choosing a statistical distribution to represent changes in the market factors and engaging in pseudo-random sampling from that distribution are simply alien to most people.

Reliability of Results. All the methods rely on historical data. Historical simulation is unique, however, in relying so directly on historical data. A danger is that the price and rate changes in the last 100 (or 500 or 1,000) days might not be typical. For example, if by chance the last 100 days were a period of low volatility in market rates and prices, the VAR computed through historical simulation will understate the risk in the portfolio. Alternatively, if by chance the U.S. dollar price of, for example, the Mexican peso rose steadily over the last 100 days and the U.S. dollar price of a peso fell on relatively few days, the VAR computed through historical simulation will come to the unrealistic conclusion that long positions in the Mexican peso involve little risk of loss. Moreover, one cannot be confident that errors of this sort will "average out." Traders will know whether the actual price changes over the last 100 days were typical and, therefore, will know for which positions the VAR is underestimated and for which it is overestimated. If VAR is used to set risk or position limits, the traders can exploit their knowledge of the biases in the VAR system and expose the company to more risk than the risk-management committee intended.

The delta-normal and Monte Carlo methods also use historical data to estimate the parameters of distributions (for example, the delta-normal methodology relies on historical data to estimate the standard deviations and correlations of a multivariate normal distribution of changes in market factors); thus, they also are subject to the risk that the historical period used might be atypical. The problem is not as severe for them, however, because assuming a particular distribution inherently limits the possible shapes that the estimated distribution can have. For example, if one assumes that the changes in the U.S. dollar price of a Mexican peso follow a normal distribution with a mean of zero, one will predict that there is a 50 percent chance that the price of a peso will fall tomorrow even if the price has risen on each of the last 100 days. Because theoretical reasoning indicates that the probability that the price of the peso will fall tomorrow is about 50 percent, regardless of what it has done over the past 100 days, this prediction is often better than the prediction implicit in historical simulation.

The delta-normal and Monte Carlo simulation methods also share another potential problem: The assumed distributions may not adequately describe the actual distributions of the market factors. Typically, actual distributions of changes in market rates and prices have fat tails relative to the normal distribution. That is, there are more realizations far away from the mean than predicted by a normal distribution. Nonetheless, for the purpose of computing VAR using a probability of 5 percent, the normal distribution assumed in the delta-normal method appears to be a reasonable approximation.[14] An issue unique to the Monte Carlo simulation method stems from the fact that the designer of the system can choose the statistical distribution to use for the market factors. This flexibility allows the designer to make poor choices, in the sense that the chosen distribution might not adequately approximate the actual distribution of market factors.

Concerns about the reliability of the methods can be partially addressed by comparing actual changes in value to VAR amounts. This sort of validation is feasible because the VAR approach explicitly specifies the probability with which actual losses will exceed the VAR amount. Validation is performed by collecting a sample of VAR amounts and actual mark-to-market portfolio profits and losses and answering two questions: First, does the distribution of actual mark-to-market profits/losses appear similar to the distribution used to determine the VAR amount? Second, do the actual losses exceed the VAR amount with the expected frequency? A limitation of this approach to validation is that chance occurrences will almost always cause the distribution of actual portfolio profits/losses to differ somewhat from the expected distribution. Therefore, reliable inferences about the quality of VAR estimates can be made only by comparing relatively large samples of VAR amounts and actual changes in portfolio value. If validation of this sort is considered essential, a short holding period must be used in computing the VAR amounts because many years will be needed to collect a large sample of monthly or quarterly VAR amounts and portfolio profits/losses.

Flexibility in Incorporating Alternative Assumptions. In some situations, the risk manager will have reason to think that the historical standard deviations and/or correlations are not reasonable estimates of future ones. For example, in the period immediately prior to the collapse of the Thai baht in July 1997, the historical correlation between changes in the USD/THB and USD/JPY exchange rates was very high, but a risk manager might have suspected that the baht might collapse and, therefore, that the correlation would be much lower in the future. How easily could the three methods have been used to calculate the VAR in this situation? Historical simulation is directly tied to the historical changes in the basic market factors. As a result, "what-if" analyses do not fit naturally within this framework.[15] In contrast, carrying out what-if analysis is easy in the delta-normal and Monte Carlo simulation methods. The user may override the parameter estimates based on historical data and use any consistent set of parameters the user chooses.

Stress Testing

If a probability of 5 percent and a holding period of one day are used in computing the VAR, the analyst should expect to suffer a loss exceeding the VAR on one (business) day out of 20, or about once a month. A level of loss that will be exceeded about once a month is reasonably termed a "normal" loss. But the next question is: When the VAR is exceeded, just how large might the loss be?

Stress testing attempts to answer this question by focusing on the losses that exceed the VAR amount. Stress testing is the rubric for a set of scenario analyses used to investigate the effects of extreme market conditions. To the extent that the effects are deemed unacceptable, the portfolio or risk-management strategy needs to be revised. There is no standard way to carry out stress testing and no standard set of scenarios to consider. Rather, the process depends crucially on the judgment and experience of the risk manager.

Stress testing often begins with a set of hypothetical extreme market scenarios. These scenarios might be created from statistical characterizations of extreme scenarios, such as assumed 5 or 10 standard deviation moves in market rates or prices. Or they might come from actual extreme events. For example, the scenarios might be based on the changes in U.S. dollar interest rates and bond prices experienced during the winter and spring of 1994 or the dramatic changes in the exchange rates of several Asian countries during the summer and fall of 1997. Alternatively, the scenarios might be created by imagining a few sudden surprises and thinking through their implications for the markets. For example, how would the unanticipated failure of a major dealer or hedge fund affect prices and liquidity in the swap markets? What would be the effect on the Korean won and the Japanese yen of the North Koreans crossing the 38th parallel?

What would be the effect of such an incident on the U.S. and Japanese equity markets? In developing these scenarios, an important aspect is to think through implications for *all* markets. An event sufficiently significant to have a sudden, major impact on the USD/JPY exchange rate would almost certainly affect other exchange rates and would likely affect interest rates in many currencies.

Companies whose risk-management strategies depend on dynamic hedging or the ability to frequently adjust or rebalance their portfolios also need to consider the impact of major surprises on market liquidity. For example, executing transactions at reasonable bid–ask spreads may be difficult or impossible during periods of market stress.

In addition, companies that use futures contracts to hedge relatively illiquid assets or financial contracts must consider the funding needs of the futures contracts. Gains or losses on futures are received or paid immediately, whereas gains or losses on other instruments are often not received or paid until the positions are closed out. As a result, even a well-hedged position that combines futures contracts with other instruments can lead to timing mismatches between when funds are required and when they are received.

The phrase "stress testing" is also used to describe scenario analyses examining the effects of violations of the assumptions underlying the VAR calculations. For example, immediately prior to the collapse of the Thai baht in July 1997, all three VAR methodologies would have indicated that from the perspective of a U.S. dollar investor, a long position in bahts combined with a short position in Japanese yen had a very low VAR. The low VAR would have been a result of the historically high correlation between the USD/THB and USD/JPY exchange rates and the fact that all three VAR methodologies rely on historical data. Yet, in July 1997, that long position in bahts would have suffered a large loss because the historical correlations no longer held. This risk could have been evaluated either by changing the correlation used as input in calculating the VAR or by examining directly the impact on portfolio value of a Thai baht fall relative to the dollar without a corresponding change in the yen. Regardless, the key input to this process is the risk manager's judgment about whether a scenario is worth considering.

Alternatives to VAR

VAR summarizes the information in the probability distribution of possible changes in portfolio value in a particular way (i.e., simply by reporting the loss that is exceeded with a probability of x percent). This measure will not be appropriate for all entities. Two alternative market risk measures are sensitivity analysis and cash flow at risk (CFAR). Sensitivity analysis is arguably less sophisticated than VAR. Companies or institutions with exposures to only a few market factors may find sensitivity analysis preferable because for them, the benefits of VAR do not justify the difficulty of mastering the approach and implementing a system to compute VAR. CFAR may be considered more sophisticated than VAR because, although CFAR is similar in concept to VAR, estimating the risk of loss in cash flows can be more complicated than estimating value at risk.

Sensitivity Analysis. The approach in sensitivity analysis is to imagine hypothetical changes in the value of each market factor, use pricing models to compute the value of the portfolio given the new value of the market factor, and then determine the change in portfolio value resulting from the change in the market factor. For example, sensitivity analysis might indicate that if the U.S. dollar price of a British pound increases by 1 percent, the value of the portfolio will decrease by USD200,000 and if the U.S. dollar price of a British pound decreases by 1 percent, the value of the portfolio will increase by USD240,000. Such computations are typically performed and reported for a range of increases and decreases that covers the range of likely exchange rate changes. Similar computations would be reported for other relevant market factors, such as interest rates.

When combined with knowledge of the magnitudes of likely exchange rate or interest rate changes, these sorts of computations provide a good picture of the risks of portfolios with exposures to only a few market factors. In fact, they compose the most basic risk-management information and are closely related to option deltas. Their principal limitation stems from the fact that a sensitivity analysis report for a portfolio with exposures to many different market factors can easily contain hundreds or thousands of numbers. In the absence of an approach like VAR, a risk manager or senior manager charged with oversight of trading and risk-management activities will find it difficult (or impossible) to read and review sensitivity analysis reports for portfolios with exposures to many different market factors and get a sense of *aggregate* portfolio risk from the information.

Cash Flow at Risk. CFAR is similar to VAR. It differs primarily in quantifying the potential loss in cash flows rather than mark-to-market values. CFAR is popular with some companies, particularly nonfinancial corporations, that are concerned more with managing the risks inherent in operating cash flows than in mark-to-market values.

CFAR is typically estimated through Monte Carlo simulation. There are important differences, however, between the use of Monte Carlo simulation to estimate VAR and its use to estimate CFAR. First, when the focus is on cash flows (not changes in mark-to-market values), hypothetical market factors must be combined with the terms of cash and derivative instruments to compute the hypothetical distribution of changes in quarterly or annual cash flows. Second, *all* future cash flows must be included in the calculation, not simply those relating to cash and derivative instruments. This comprehensiveness is essential if the goal of risk measurement is to assess the impact of derivatives and other financial transactions on companies' total cash flows, including *operating* cash flows. As a result, the factors to be included in the simulation are not simply the basic financial market factors included in VAR calculations; they include any factors that affect operating cash flows. Changes in customer demand, the outcomes of research and development programs, and competitors' pricing decisions are a few examples. Third, the time horizon is typically much longer in CFAR simulation because of the longer planning cycles of nonfinancial companies. For example, in CFAR, values of the underlying market factors might be simulated for the next 20 quarters, whereas the time horizon in VAR computations is often as short as one day. Finally, the primary objective in using CFAR is often to facilitate internal planning rather than to oversee and control firm risk.

A serious drawback is that successful design and implementation of a CFAR measurement system requires a high degree of knowledge and judgment.[16] First, the designer must determine the important operating factors and how they affect operating cash flows in order to develop a model of the company's operating cash flows. This step alone may be a major undertaking. Next, this model of operating cash flows must be integrated with a model of financial market factors. Then, the user must select the statistical distribution from which the hypothetical values of the factors (both operating and financial) are drawn and select or estimate the parameters of that distribution. This step can be particularly difficult in the case of the operating factors. In contrast to the financial market factors, data on actual past changes in operating risk factors may not be available to guide the choice of distribution. Finally, the user must carry out the computations. Somewhat offsetting these difficulties is that the model of the financial market factors can be relatively crude because there is no point in refining it to be more precise than the model of the operating cash flows. Nonetheless, building a CFAR measurement system is likely to be a major undertaking.

Conclusion: Is VAR for You?

VAR is not a panacea. It is a single, summary statistical measure of normal market risk. As pointed out here, VAR summarizes the information in the probability distribution of possible changes in portfolio value in a particular way (i.e., simply by reporting the loss that is exceeded with a probability of x percent). At the level of the trading desk, it is simply one more item in the risk manager's or trader's toolkit. Traders and frontline risk managers will look at the whole panoply of Greek letter risks (the deltas, gammas, vegas, etc.) and may look at the portfolio's exposures to other factors, such as changes in correlations. The only environment in which VAR numbers will be used alone is at the level of oversight by senior managers or regulators. Even at this level, results of scenario analysis, stress tests, and other information about the positions will often supplement the VAR numbers. Clearly, VAR estimates do not contain all the information about market risks that one would like to have.

Moreover, VAR is based on a range of assumptions, few (or none) of which will be satisfied exactly. Perhaps most importantly, it is an estimate of risk, often based on historical data, that relies on the idea that the future will be like the past. For this reason and others, it has been increasingly criticized (for an example, see Taleb 1996). Given these criticisms, a person could quite reasonably wonder whether there is sufficient information in VAR estimates to justify the costs of producing them. But this skepticism needs to be tempered by a consideration of the alternatives—which may be as unappealing as showing your boss a table with hundreds of market factor sensitivities.

Notes

1. The option delta measures the sensitivity of the option value to changes in the value of the underlying asset and is defined as $\Delta = \partial V/\partial S$, where V is the value of the option and S is the price of the underlying asset. The option gamma, defined as $\Gamma = \partial\Delta/\partial S = \partial 2V/\partial S2$, measures how delta changes as the price of the underlying asset changes. See Hull (2000, Chapter 13) or Kolb (1997, Chapter 14) for more discussion of these concepts.

2. Netting the risk makes sense because gains or losses on the long position in markka will offset gains or losses on the short position in krona.
3. Your answer should not start: "The most we can lose is . . ." because the only honest way to finish this sentence is "everything."
4. As you will see in the discussion of the historical simulation method, the daily VAR using a 5 percent probability is actually USD97,230.
5. As suggested by its name, the delta-normal method assumes that the distributions of the underlying market risk factors and the portfolio value are normal. Under this assumption, the loss exceeds 1.645 times the standard deviation of portfolio value with a probability of 5 percent and exceeds 2.326 times the standard deviation of portfolio value with a probability of 1 percent. The ratio of these figures is 2.326/1.645 = 1.414.
6. In some cases, formulas are not available and the instruments' values must be computed using numerical algorithms.
7. The maturities need not be the same for every currency. The interest rates for long maturities typically will not be relevant for currencies in which no active long-term debt markets exist.
8. The delta-normal approach is also called the "parametric" approach or "variance–covariance" approach.
9. This procedure of using the May 20th market factors together with the historical changes to generate hypothetical May 21st market factors makes sense because it guarantees that the hypothetical May 21st values will be more or less centered around the May 20th values, which is reasonable because the May 20th daily VAR is a measure of the portfolio gain or loss that might occur during the next trading day. An alternative procedure of computing the hypothetical mark-to-market portfolio values using the actual levels of the market factors observed over the past 100 days frequently involves using levels of the market factors that are not close to the current values. This reasoning does not imply, however, that one must use percentage changes together with the May 20th values. Alternatives are to use logarithmic changes or "absolute" changes. By using percentage changes, we are implicitly assuming that the statistical distribution of percentage changes in the market factors does not depend on their levels.
10. For example, to measure the risk of changes in credit spreads between the yields on corporate and government bonds, a VAR system has to include as market factors the yields on both government and corporate bonds. If the market factors include only government bond yields, the system will implicitly assume that changes in the prices of corporate bonds are perfectly correlated with changes in the prices of government bonds. Thus, the system will be unable to capture the risk of changes in credit or yield spreads.
11. The standard deviation is, of course, simply the square root of the variance.
12. The natural interpretations of its parameters (means, standard deviations, and correlations) and the ease with which these parameters can be estimated weigh in favor of the multivariate normal distribution.
13. The pricing models need not be perfect, however, because VAR focuses on changes in value. If the error in the pricing model is reasonably stable, in the sense that the error in today's price is about the same as the error in tomorrow's, then changes in value computed using the pricing model will be correct even though the level of the prices is not.
14. A good discussion of this issue can be found in J.P. Morgan (1995).
15. Special procedures must be used if such scenarios are to be incorporated in historical simulation. For example, Duffie and Pan (1997) suggested that alternative assumptions about the standard deviations of a market factor can be incorporated by subtracting the mean change in the market factor from the vector of changes and then multiplying the result by a constant to rescale the changes in the market factor.
16. Off-the-shelf software is currently not available, and it may never be available, because CFAR systems typically include operating cash flows, the characteristics of which are company specific and difficult to incorporate in any general system. At least one major derivatives dealer has been willing, however, to provide current and potential future customers with the framework of a CFAR system, the simulation engines, and assistance in implementing the system (Hayt and Song 1995).

References

Basle Committee on Banking Supervision. 1996. "Amendment to the Capital Accord to Incorporate Market Risks." Basle, Switzerland: Basle Committee on Banking Supervision (January).

Duffie, D., and J. Pan. 1997. "An Overview of Value at Risk." *Journal of Derivatives*, vol. 4, no. 3 (Spring):7–49.

Fong, H.G., and K. Lin. 1999. "A New Analytical Approach to VAR." *Journal of Portfolio Management*, 25 Anniversary issue (May):88–97.

Hayt, G., and S. Song. 1995. "Handle with Sensitivity." *Risk*, vol. 8, no. 9 (September):94–99.

Hull, John C. 2000. *Options, Futures, and Other Derivatives*. 4th ed. Upper Saddle River, NJ: Prentice-Hall.

J.P. Morgan. 1994. *RiskMetrics Technical Document*. New York: Morgan Guaranty Trust Company (October).

———. 1995. *RiskMetrics Technical Document*. 3rd ed. New York: Morgan Guaranty Trust Company (May).

———. 1996. *RiskMetrics Monitor*. New York: Morgan Guaranty Trust Company (First Quarter).

J.P. Morgan/Reuters. 1996. *RiskMetrics Monitor*. New York: Morgan Guaranty Trust Company (Third Quarter).

Kolb, Robert W. 1997. *Futures, Options, and Swaps*. Malden, MA: Blackwell Publishers.

Taleb, Nassim. 1996. *Dynamic Hedging: Managing Vanilla and Exotic Options*. New York: Wiley.

U.S. Securities and Exchange Commission. 1997. *Disclosure of Accounting Policies for Derivative Financial Instruments and Derivative Commodity Instruments and Disclosure of Quantitative and Qualitative Information about Market Risk Inherent in Derivative Financial Instruments, Other Financial Instruments, and Derivative Commodity Instruments*. Release Nos. 33-7386; 34-38223; IC-22487; FR-48; International Series No. 1047; File No. S7-35-95. Washington, DC: SEC (January 31).

Appendix A. Theory Underlying Mapping the Forward Contract into the Three Standardized Positions

As discussed in the text, the idea behind risk mapping is to replace the actual portfolio with a portfolio of standardized positions that has the same exposures or sensitivities to the basic market factors. If the market factors are chosen so that they capture all (or most) of the risks to which the portfolio is exposed, then the risk of the portfolio of standardized positions will be (approximately) the same as the risk of the actual portfolio. We show that the forward contract used in the examples does, in fact, have the same exposures or sensitivities as the portfolio of the three standardized positions to which it was mapped. The purpose is to justify the mapping we performed.

In this context, matching exposures means matching partial derivatives, or deltas. We illustrate this point by using first-order Taylor series approximations to represent the changes in the values of both the forward contract and the portfolio of the three standardized positions in terms of changes in the three market factors, and we choose the standardized positions so that the coefficients of the two Taylor series approximations are the same. If the coefficients of the approximations are the same, then (up to the approximation) the two portfolios respond identically to changes in the market factors.

First, we consider the forward contract. Let the mark-to-market value of the forward contract, V_F, be denoted by

$$V_F \equiv S \times \left[\frac{\text{GBP10 million}}{1 + r_{GBP}(91/360)}\right] - \frac{\text{USD15 million}}{1 + r_{USD}(91/360)},$$

where S is the spot exchange rate and r is the interest rate.

Using a Taylor series, the change in V_F can be approximated as follows:

$$\begin{aligned}\text{Change in } V_F &\approx \frac{\partial V_F}{\partial r_{USD}}(\text{Change in } r_{USD}) + \frac{\partial V_F}{\partial r_{GBP}}(\text{Change in } r_{GBP}) + \frac{\partial V_F}{\partial S}(\text{Change in } S)\\ &= \Delta^F_{r_{USD}}(\text{Change in } r_{USD}) + \Delta^F_{r_{GBP}}(\text{Change in } r_{GBP}) + \Delta^F_S(\text{Change in } S),\end{aligned}$$

where $\Delta^F_{r_{USD}}$, $\Delta^F_{r_{GBP}}$, and Δ^F_S are the deltas of the forward contract with respect to the three market factors.

Next, we write down a similar Taylor series approximation of the change in the value of the portfolio of standardized positions (i.e., change in V) and show that if the standardized positions are chosen appropriately, the coefficients of the two approximations are identical. In that case, a change in $V \approx$ a change in V_F, which implies that (up to the approximation) the portfolio of standardized positions has the same sensitivities to the market factors as the forward contract.

Let $V \equiv X_1 + X_2 + X_3$ represent the value of the portfolio of standardized positions. If each of the X's depends on only one market factor, then the change in V can be approximated as follows:

$$\begin{aligned}\text{Change in } V &\approx \frac{\partial X_1}{\partial r_{USD}}(\text{Change in } r_{USD}) + \frac{\partial X_2}{\partial r_{GBP}}(\text{Change in } r_{GBP}) + \frac{\partial X_3}{\partial S}(\text{Change in } S)\\ &= \Delta_{r_{USD}}(\text{Change in } r_{USD}) + \Delta_{r_{GBP}}(\text{Change in } r_{GBP}) + \Delta_S(\text{Change in } S),\end{aligned}$$

where $\Delta_{r_{USD}}$, $\Delta_{r_{GBP}}$, and Δ_S are the deltas of the standardized positions with respect to the three market factors.

We need to choose X_1, X_2, and X_3 so that each depends on only one market factor and the two Taylor series approximations are identical. This task amounts to choosing them so that $\partial X_1/\partial r_{USD} = \partial V_F/\partial r_{USD}$, $\partial X_2/\partial r_{GBP} = \partial V_F/\partial r_{GBP}$, and $\partial X_3/\partial S = \partial V_F/\partial S$ or, equivalently, to matching the corresponding deltas. The choice that works is

$$\begin{aligned}X_1 &= \left.\frac{\text{USD15 million}}{1 + r_{USD}(91/360)}\right|_{r_{USD} = 0.0546875}\\ &= -\text{USD14,795,000};\\ X_2 &= \left.\frac{(1.5355\ \text{USD/GBP}) \times \text{GBP15 million}}{1 + r_{GBP}(91/360)}\right|_{r_{GBP} = 0.060625}\\ &= \text{USD15,123,242};\\ X_3 &= \left.(S\ \text{USD/GBP})\frac{\text{GBP15 million}}{1 + 0.060625(91/360)}\right|_{S = 1.5355}\\ &= \text{USD15,123,242}.\end{aligned}$$

These positions are the three standardized positions that represent the risk mapping of the forward contract as described in the body of the article.

Appendix B. Mapping Options and Other Positions

Although the approach for mapping options can be more complicated, it is no different in concept from the approach for mapping the forward contract discussed in Appendix A. As with the forward contract, we replace the actual portfolio with a portfolio of standardized positions that has the same exposures or sensitivities (i.e., option deltas) to the basic market factors. We can proceed in two ways, both of which result in the identical portfolio of standardized positions. We illustrate both of them by considering an OTC option on a 10-year British government bond, or gilt. We assume that in the VAR system the 20 semiannual cash flows of a 10-year gilt will be mapped onto the three-month and six-month and 1-, 2-, 3-, 4-, 5-, 7-, and 10-year British-pound-denominated zero-coupon bonds, plus a position in "spot" pounds, for a total of 10 standardized positions.

Ordinarily, one would say that the underlying asset of the option is the 10-year gilt. Taking this point of view, the first approach maps the option to a delta-equivalent position in the 10-year gilt. To understand this process, recall that the option delta indicates the price change in the option resulting from a one-unit change in the price of the underlying asset. For example, an option with a delta of 1/2 changes in value by GBP1/2 when the gilt changes in value by GBP1. That is, the option has the risk of $\Delta = 1/2$ gilts. Using this idea, the first step in this approach is to map the option to a delta-equivalent position of Δ gilts. But gilts themselves usually are not a standardized position in most risk-measurement systems, so a second step is necessary. This step maps the position of Δ gilts into the standardized positions in the risk-measurement system. The mapping is performed in the same way any other bond or forward contract would be mapped to standardized positions.

The alternative approach is to map the option to the standardized positions in a single step. If we take the perspective of a pound-based investor, we can interpret the option on the gilt as an option on a portfolio of the nine pound-denominated zero-coupon bonds and think of the option as having nine underlying assets and nine deltas, one for each underlying asset. The U.S. dollar price of the gilt also depends, however, on the USD/GBP exchange rate, which gives it a tenth underlying asset and a tenth delta. Thus, from the perspective of a dollar-based investor, there are 10 underlying assets—the 9 pound-denominated zero-coupon bonds and the USD/GBP exchange rate, and for each, we can define a delta.

Letting V denote the dollar value of the gilt and P_n represent the British pound price of the nth pound-denominated zero-coupon bond, we have for the first nine deltas

$$\Delta_n = \frac{\partial V}{\partial P_n} \text{ for } n = 1, \ldots, 9,$$

and for the 10th delta, the partial derivative with respect to the spot exchange rate, we have

$$\Delta_{10} = \frac{\partial V}{\partial S}.$$

Using these 10 deltas, we map the option to the 10 standardized positions in a single step by finding a portfolio of the standardized positions that has the same deltas with respect to the market factors.

A limitation of the delta-normal approach is precisely that it is based on mapping instruments to their delta-equivalent positions (i.e., it amounts to finding a portfolio of standardized positions that have the same deltas as the options). Using delta as a risk measure is valid for small changes in the underlying market factors, but if instruments' deltas change as market factors change, delta does not properly measure the changes in value resulting from large changes in factors. Thus, to the extent that instruments' deltas change as market factors change, mapping based on delta will introduce errors in the approach.

We can think about these errors in terms of the option gammas (Γ's), which supplement the deltas by measuring how the deltas change as the price of the underlying asset, currency, or commodity changes. The option gamma is defined as the partial derivative of delta with respect to the price of the underlying asset, currency, or commodity or, equivalently, as the second partial derivative of the option price with respect to the price of the underlying asset, currency, or commodity.

Letting S denote the spot price of the underlying asset and $V(S)$ denote the option price as a function of S, the option gamma is

$$\Gamma = \frac{\partial \Delta(S)}{\partial S} = \frac{\partial^2 V(S)}{\partial S^2}.$$

Delta and gamma together can be used to predict the change in the option price resulting from a change in S by using the following formula:

$$\text{Change in option price} \approx \Delta \times \begin{pmatrix} \text{Change in} \\ \text{market factor} \end{pmatrix} + \frac{1}{2}\Gamma \times \begin{pmatrix} \text{Change in} \\ \text{market factor} \end{pmatrix}^2 .$$

In contrast, using delta alone produces

$$\text{Change in option price} \approx \Delta \times \begin{pmatrix} \text{Change in} \\ \text{market factor} \end{pmatrix}.$$

Comparing these two equations shows that the error in relying on delta is based on the magnitude and sign of gamma: When gamma is negative, the change in the option price is more adverse than that predicted using delta alone; when gamma is positive, the change in the option price is more favorable than that predicted using delta alone.

The significance of this aspect for VAR measures is that the delta-normal method typically uses delta alone in measuring the risk of options; therefore, it somewhat understates the risk of positions with negative gammas. The effect will be small for VAR computations for short holding periods because the change in the spot price of the underlying asset is typically small for short periods and the term $1/2\Gamma$(Change in price of underlying instrument)^2 is small. The understatement of the risk of negative gamma portfolios can be significant, however, when VAR measures are computed for long holding periods. **Figure B1** illustrates an extreme version of the problem.

Figure B1. Risky Portfolio That Has Delta = 0

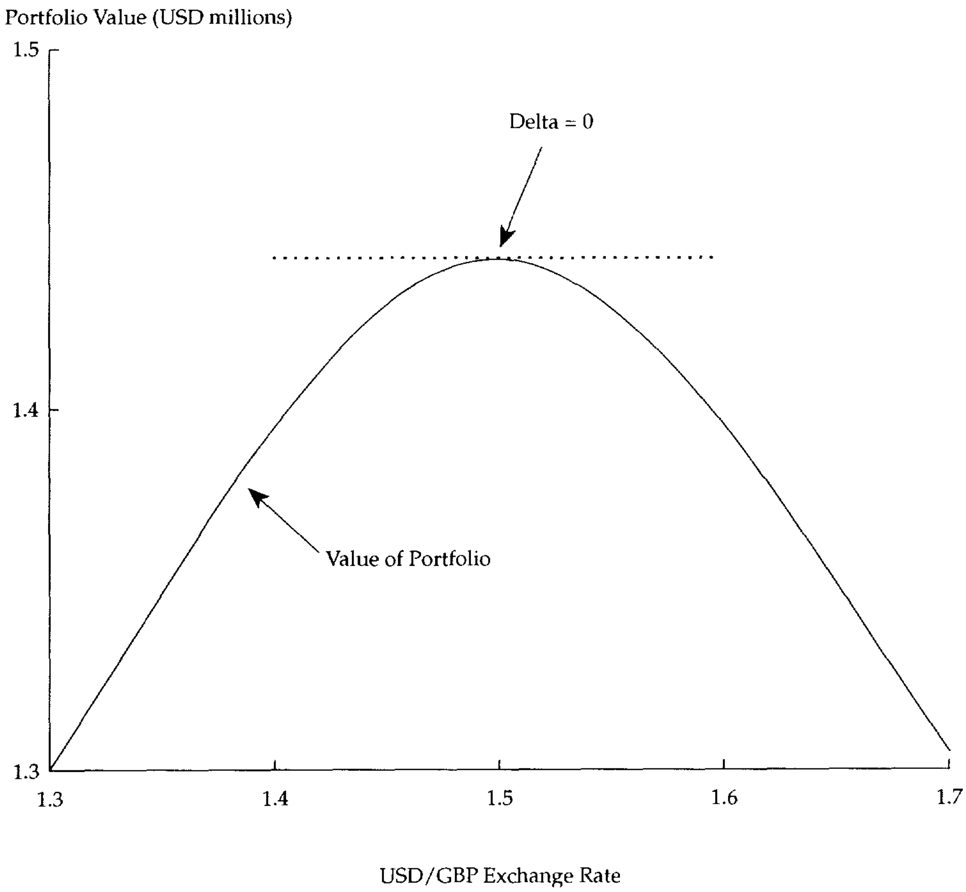

Figure B1 shows the value of a portfolio exposed to the USD/GBP exchange rate as a function of the current value of the exchange rate. At the current exchange rate of 1.5 USD/GBP, the portfolio delta is zero and the delta-normal approach would map the portfolio to a position of zero British pounds and thus estimate the VAR to be zero. The portfolio is clearly risky, however, in that a change in the USD/GBP exchange rate in either direction will lead to a loss. The downward curvature in the portfolio value function corresponds to a negative gamma and illustrates that the delta-normal approach can underestimate the risk of option portfolios with negative gamma.

This problem can be overcome by using Monte Carlo simulation. As alternatives, various delta-gamma approaches have been proposed to deal with the problem (J.P. Morgan 1996; J.P. Morgan/Reuters 1996), although none are fully satisfactory. Recently, Fong and Lin (1999) proposed a promising new analytical approach that allows for easy calculation of the VAR of option portfolios.