

Risk Measurement: An Introduction to Value at Risk

Thomas J. Linsmeier and Neil D. Pearson*

University of Illinois at Urbana-Champaign

July 1996

Abstract

This paper is a self-contained introduction to the concept and methodology of “value at risk,” which is a new tool for measuring an entity’s exposure to market risk. We explain the concept of value at risk, and then describe in detail the three methods for computing it: historical simulation; the variance-covariance method; and Monte Carlo or stochastic simulation. We then discuss the advantages and disadvantages of the three methods for computing value at risk. Finally, we briefly describe some alternative measures of market risk.

© Thomas J. Linsmeier and Neil D. Pearson

* Department of Accountancy and Department of Finance, respectively. Linsmeier may be reached at 217 244 6153 (voice), 217 244 0902 (fax), and linsmeie@uiuc.edu; Pearson may be reached at 217 244 0490 (voice), 217 244 9867 (fax), and pearson2@uiuc.edu. We are solely responsible for any errors.

A DIFFICULT QUESTION

You are responsible for managing your company's foreign exchange positions. Your boss, or your boss's boss, has been reading about derivatives 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 positions, but this isn't likely to be helpful unless there are only a handful. Even then, it helps only if your superiors understand all of the positions and instruments, and the risks inherent in each. Or you could talk about the portfolio's sensitivities, i.e. how much the value of the portfolio changes when various underlying market rates or prices change, and perhaps option delta's and gamma's.¹ However, you are unlikely to win favor with your superiors by putting them to sleep. Even if you are confident in your ability to explain these in English, you still have no natural way to net the risk of your short position in Deutsche marks against the long position in Dutch guilders. (It makes sense to do this because gains or losses on the short position in marks will be almost perfectly offset by gains or losses on the long position in guilders.) You could simply assure your superiors that you never speculate but rather use derivatives only to hedge, but they understand that this statement is vacuous. They know 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?

Perhaps the best answer starts: "The value at risk is"²

How did you get into a position where the best answer involves a concept your superiors might never have heard of, let alone understand? This doesn't seem like a good strategy for getting promoted.

The modern era of risk measurement for foreign exchange positions began in 1973. That year saw both the collapse of the Bretton Woods system of fixed exchange rates and the publication of the Black-Scholes option pricing formula. The collapse of the Bretton Woods system and the rapid transition to a system of more or less freely floating exchange rates among many of the major trading countries provided the impetus for the measurement and management of foreign exchange risk, while the ideas underlying the Black-Scholes formula provided the conceptual framework and basic tools for risk measurement and management.

The years since 1973 have witnessed both tremendous volatility in exchange rates and a proliferation of derivative instruments useful for managing the risks of changes in the prices of foreign currencies and interest rates. Modern derivative instruments such as forwards, futures, swaps, and options facilitate the management of exchange and interest rate volatility. They can be used to offset the risks in existing instruments, positions, and portfolios because their cash flows and values change with changes in interest rates and foreign currency prices. Among other things, they can be used to make offsetting bets to "cancel out" the risks in a portfolio. Derivative instruments are ideal for this purpose, because many of them can be traded quickly, easily, and with low transactions costs, while others can be tailored to customers' needs. Unfortunately,

¹ Option delta's and gamma's are defined in Appendix A.

² Your answer doesn't start: "The most we can lose is ..." because the only honest way to finish this sentence is "everything." It is possible, though unlikely, that all or most relevant exchange rates could move against you by large amounts overnight, leading to losses in all or most currencies in which you have positions.

instruments which are ideal for making offsetting bets also are ideal for making purely speculative bets: offsetting and purely speculative bets are distinguished only by the composition of the rest of the portfolio.

The proliferation of derivative instruments has been accompanied by increased trading of cash instruments and securities, and has been coincident with growth in foreign trade and increasing international financial linkages among companies. As a result of these trends, many companies have portfolios which include large numbers of cash and derivative instruments. Due to the sheer numbers and complexity (of some) of these cash and derivative instruments, the magnitudes of the risks in companies' portfolios often are not obvious. This has led to a demand for portfolio level quantitative measures of market risk such as "value at risk." The flexibility of derivative instruments and the ease with which both cash and derivative instruments can be traded and retraded to alter companies' risks also has created a demand for a portfolio level summary risk measure that can be reported to the senior managers charged with the oversight of risk management and trading operations.

The ideas underlying option pricing provide the foundation for the measurement and management of the volatility of market rates and prices. The Black-Scholes model and its variants had the effect of disseminating probabilistic and statistical tools throughout financial institutions and companies' treasury groups. These tools permit quantification and measurement of the volatility in foreign currency prices and interest rates. They are the foundation of value at risk and risk measurement systems. Variants of the Black-Scholes model, known as the Black and Garman-Kohlhagen models, are widely used for pricing options on foreign currencies and foreign currency futures. Most other pricing models are also direct descendants of the Black-Scholes model. Even the pricing of simpler instruments such as currency and interest rate swaps is based on the "no-arbitrage" framework underlying the Black-Scholes model. Partial derivatives of various pricing formulas provide the basic risk measures. These basic risk measures are discussed in the first appendix to this chapter.

The concept and use of value at risk is recent. Value at risk was first used by major financial firms in the late 1980's to measure the risks of their trading portfolios. Since that time period, the use of value at risk has exploded. Currently value at risk is used by most major derivatives dealers to measure and manage market risk. In the 1994 follow-up to the survey in the Group of Thirty's 1993 global derivatives project, 43% of dealers reported that they were using some form of value at risk and 37% indicated that they planned to use value at risk by the end of 1995. J.P. Morgan's attempt to establish a market standard through its release of its RiskMetrics™ system in October 1994 provided a tremendous impetus to the growth in the use of value at risk. Value at risk is increasingly being used by smaller financial institutions, non-financial corporations, and institutional investors. The 1995 Wharton/CIBC Wood Gundy *Survey of derivatives usage among US non-financial firms* reports that 29% of respondents use value at risk for evaluating the risks of derivatives transactions. A 1995 *Institutional Investor* survey found that 32% of firms use value at risk as a measure of market risk, and 60% of pension funds responding to a survey by the New York University Stern School of Business reported using value at risk.

Regulators also have become interested in value at risk. In April 1995, the Basle Committee on Banking Supervision proposed allowing banks to calculate their capital requirements for market risk with their own value at risk models, using certain parameters provided by the committee. In June 1995, the US Federal Reserve proposed a "precommitment" approach which would allow banks to use their own internal value at risk models to calculate capital requirements for market

risk, with penalties to be imposed in the event that losses exceed the capital requirement. In December 1995, the US Securities and Exchange Commission released for comment a proposed rule for corporate risk disclosure which listed value at risk as one of three possible market risk disclosure measures. The European Union's Capital Adequacy Directive which came into effect in 1996 allows value at risk models to be used to calculate capital requirements for foreign exchange positions, and a decision has been made to move toward allowing value at risk to compute capital requirements for other market risks.

SO WHAT IS VALUE AT RISK, ANYWAY?

Value at risk is a single, summary, statistical measure of possible portfolio losses. Specifically, value at risk is a measure of losses due to "normal" market movements. Losses greater than the value at risk are suffered only with a specified small probability. Subject to the simplifying assumptions used in its calculation, value at risk 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 value at risk is straightforward to understand. It is simply a way to describe the magnitude of the likely losses on the portfolio.

To understand the concept of value at risk, consider a simple example involving an FX forward contract entered into by a U.S. company at some point in the past. Suppose that the current date is 20 May 1996, and the forward contract has 91 days remaining until the delivery date of 19 August. The 3-month US dollar (USD) and British pound (GBP) interest rates are $r_{\text{USD}} = 5.469\%$ and $r_{\text{GBP}} = 6.063\%$, respectively, and the spot exchange rate is 1.5335 \$/£. On the delivery date the U.S. company will deliver \$15 million and receive £10 million. The US dollar mark-to-market value of the forward contract can be computed using the interest and exchange rates prevailing on 20 May. Specifically,

$$\begin{aligned} \text{USD mark - to - market value} &= \left[(\text{exchange rate in USD / GBP}) \times \frac{\text{GBP 10 million}}{1 + r_{\text{GBP}}(91/360)} \right] - \frac{\text{USD 15 million}}{1 + r_{\text{USD}}(91/360)} \\ &= \left[(1.5335 \text{ USD / GBP}) \times \frac{\text{GBP 10 million}}{1 + 0.06063(91/360)} \right] - \frac{\text{USD 15 million}}{1 + 0.05469(91/360)} \\ &= \text{USD 327,771.} \end{aligned}$$

In this calculation we use that fact that one leg of the forward contract is equivalent to a pound-denominated 91-day zero coupon bond and the other leg is equivalent to a dollar-denominated 91-day zero coupon bond.

On the next day, 21 May, it is likely that interest rates, exchange rates, and thus the value of the forward contract have all changed. Suppose that the distribution of possible one day changes in the value of the forward contract is that shown in Figure 1. The figure indicates that the probability that the loss will exceed \$130,000 is two percent, the probability that the loss will be between \$110,000 and \$130,000 is one percent, and the probability that the loss will be between \$90,000 and \$110,000 is two percent. Summing these probabilities, there is a five percent

probability that the loss will exceed approximately \$90,000.³ If we deem a loss that is suffered less than 5 percent of the time to be a loss due to unusual or “abnormal” market movements, then \$90,000 divides the losses due to “abnormal” market movements from the “normal” ones. If we use this 5 percent probability as the cutoff to define a loss due to normal market movements, then \$90,000 is the (approximate) value at risk.

The probability used as the cutoff need not be 5 percent, but rather is chosen by either the user or the provider of the value at risk number: perhaps the risk manager, risk management committee, or designer of the system used to compute the value at risk. If instead the probability were chosen to be two percent, the value at risk would be \$130,000, because the loss is predicted to exceed \$130,000 only two percent of the time.

Also, implicit in this discussion has been a choice of holding period: Figure 1 displays the distribution of *daily* profits and losses. One also could construct a similar distribution of 5-day, or 10-day, profits and losses, or perhaps even use a longer time horizon. Since 5 or 10-day profits and losses typically are larger than 1-day profits and losses, the distributions would be more disperse or spread out, and the loss that is exceeded only 5 (or 2) percent of the time would be larger. Therefore the value at risk would be larger.

Now that we’ve seen an example of value at risk, we are ready for the definition. Using a probability of x percent and a holding period of t days, an entity’s *value at risk* is the loss that is expected to be exceeded with a probability of only x percent during the next t -day holding period. Loosely, it is the loss that is expected to be exceeded during x percent of the t -day holding periods. Typical values for the probability x are 1, 2.5, and 5 percent, while common holding periods are 1, 2, and 10 (business) days, and 1 month. The theory provides little guidance about the choice of x . It is determined primarily by how the designer and/or user of the risk management system wants to interpret the value at risk number: is an “abnormal” loss one that occurs with a probability of 1 percent, or 5 percent? For example, JP Morgan’s RiskMetrics™ system uses 5 percent, while Mobil Oil’s 1994 annual report indicates that it uses 0.3 percent. The parameter t is determined by the entity’s horizon. Those which actively trade their portfolios, such as financial firms, typically use 1 day, while institutional investors and non-financial corporations may use longer holding periods. A value at risk number applies to the current portfolio, so a (sometimes implicit) assumption underlying the computation is that the current portfolio will remain unchanged throughout the holding period. This may not be reasonable, particularly for long holding periods.

In interpreting value at risk numbers, it is crucial to keep in mind the probability x and holding period t . Without them, value at risk numbers are meaningless. For example, two companies holding identical portfolios will come up with different value at risk estimates if they make different choices of x and t . Obviously, the loss that is suffered with a probability of only 1 percent is larger than the loss that is suffered with a probability of 5 percent. Under the assumptions used in some value at risk systems, it is 1.41 times as large.⁴ The choice of holding

³ As we will see in the discussion of the historical simulation method, the daily value at risk using a 5% probability is actually \$97,230.

⁴ The variance-covariance 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 is $1.414=2.326/1.645$.

period can have an even larger impact, for the value at risk computed using a t -day holding period is approximately $\sqrt{t}$ times as large as the value at risk using a one day holding period. Absent appropriate adjustments for these factors, value at risk numbers are not comparable across entities.

Despite its advantages, value at risk is not a panacea. It is a single, summary, statistical measure of normal market risk. At the level of the trading desk, it is just one more item in the risk manager's or trader's toolkit. The traders and front-line risk managers will look at the whole panoply of Greek letter risks, i.e. the delta's, gamma's, vega's, et cetera, and may look at the portfolio's exposures to other factors such as changes in correlations. In many cases they will go beyond value at risk and use simulation techniques to generate the entire distribution of possible outcomes, and will supplement this with detailed analyses of specific scenarios and "stress tests." The only environment in which value at risk numbers will be used alone is at the level of oversight by senior management. Even at this level, the value at risk numbers often will be supplemented by the results of scenario analyses, stress tests, and other information about the positions.

In the balance of this chapter we describe the three main methods for computing value at risk numbers: historical simulation, the variance-covariance or analytic method, and Monte Carlo or stochastic simulation. We then consider the advantages and disadvantages of the three methods, how they can be supplemented with "stress testing," and a brief discussion of some of the alternatives to value at risk. Appendices to the chapter review option delta's and gamma's and explain the concept of "risk mapping" which is used in the variance-covariance method. First, however, we need to discuss a fundamental idea which underlies value at risk computations.

FUNDAMENTALS: IDENTIFYING THE IMPORTANT MARKET FACTORS

In order to compute value at risk (or any other quantitative measure of market risk), we need to identify the basic market rates and prices that affect the value of the portfolio. These basic market rates and prices are the "market factors." It is necessary to identify a limited number of basic market factors simply because otherwise the complexity of trying to come up with a portfolio level quantitative measure of market risk explodes. Even if we restrict our attention to simple instruments such as forward contracts, an almost countless number of different contracts can exist, because virtually any forward price and delivery date are possible. The market risk factors inherent in most other instruments such as swaps, loans (often with embedded options), options, and exotic options of course are ever more complicated. Thus, expressing the instruments' values in terms of a limited number of basic market factors is an essential first step in making the problem manageable.

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 illustrate this using the FX forward contract we introduced above. The current date is 20 May 1996. The contract requires a US company to deliver \$15 million in 91 days. In exchange it will receive £10 million. The current US dollar market value of this forward contract depends on three basic market factors: S , the spot exchange rate expressed in dollars per pound; r_{GBP} , the 3-month pound interest rate; and r_{USD} , the 3-month dollar interest rate. To see this, we decompose the cash flows of the forward contract into the following equivalent portfolio of zero-coupon bonds:

Position	Current \$ Value of Position	Cash Flow on Delivery Date
Long position in 91 day £ denominated zero coupon bond with face value of £10 million	$S \times \frac{\text{GBP 10 million}}{1+r_{\text{GBP}}(91/360)}$	Receive £10 million
Short position in 91 day \$ denominated zero coupon bond with face value of \$15 million	$-\frac{\text{USD 15 million}}{1+r_{\text{USD}}(91/360)}$	Pay \$15 million

The decomposition yields the following formula, used above, for the current mark-to-market value (in dollars) of the position in terms of the basic market factors r_{USD} , r_{GBP} , and S :

$$\text{USD mark - to - market value} = \left[S \times \frac{\text{GBP 10 million}}{1 + r_{\text{GBP}}(91/360)} \right] - \frac{\text{USD 15 million}}{1 + r_{\text{USD}}(91/360)}.$$

Because this is an over-the-counter forward contract subject to some credit risk, the interest rates are those on 3-month interbank deposits (LIBOR) rather than the rates on government securities. Similar formulas expressing the instruments' values in terms of the basic market factors must be obtained for all of the instruments in the portfolio.⁵ Once such formulas have been obtained, a key part of the problem of quantifying market risk has been finished. The remaining steps involve determining or estimating the statistical distribution of the potential future values of the market factors, using these potential future values and the formulas to determine potential future changes in the values of the various positions that comprise the portfolio, and then aggregating across positions in order to determine the potential future changes in the value of the portfolio. Value at risk is a measure of these potential changes in the portfolio's value.

Of course, the values of most actual portfolios will depend upon 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 with a range of maturities. For example, the maturities used in the first version of JP Morgan's 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.⁶ A company with positions in most of the actively traded currencies, and a number of the minor ones, could easily have a portfolio exposed to several hundred market factors.

This dependence on only a limited number of basic market factors typically remains implicit in the historical and Monte Carlo simulation methodologies, but must be made explicit in the variance-covariance methodology. The process of making this dependence explicit is known as "risk mapping." Specifically, risk mapping involves taking the actual instruments and "mapping" them into a set of simpler, standardized positions or instruments. We describe this process when we discuss the variance-covariance method below, and in Appendix B.

⁵ In some cases formulas are not available and instruments' values must be computed using numerical algorithms.

⁶ 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 there are not active long term debt markets.

VALUE AT RISK METHODOLOGIES

Historical simulation

Historical simulation is a simple, atheoretical approach that requires relatively few assumptions about the statistical distributions of the underlying market factors. We illustrate the procedure with a simple portfolio consisting of a single instrument, the 3-month FX forward for which the distribution of hypothetical mark-to-market profits and losses was previously shown in Figure 1. In essence, the approach involves using historical changes in market rates and prices to construct a distribution of potential future portfolio profits and losses in Figure 1, and then reading off the value at risk as the loss that is exceeded only 5% 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, here days. That is, N sets of hypothetical market factors are constructed using 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. Doing this allows one to compute N hypothetical mark-to-market profits and losses on the portfolio, when compared to the current mark-to-market portfolio value. Even though the actual changes in rates and prices are used, the mark-to-market profits and losses are hypothetical because the current portfolio was not held on each of the last N periods. The use of the actual historical changes in rates and prices to compute the hypothetical profits and losses is the distinguishing feature of historical simulation, and the source of the name. Below we illustrate exactly how to do this. 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 value at risk, can then be determined.

Performing the analysis for a single instrument portfolio

We carry out the analysis as of the close of business on 20 May, 1996. Recall that the forward contract obligates a U.S. company to deliver \$15 million on the delivery date 91 days hence, and in exchange receive £10 million. We perform the analysis from the perspective of the US 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. For simplicity, we assume that the holding period is one day ($t=1$), the value at risk will be computed using a 5 percent probability ($\alpha=5\%$), and that 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 20 May is the 100th business day of 1996, the most recent 100 business days start on 2 January 1996.

Historical simulation can be described in terms of 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: they are the 3-month pound interest rate, the 3-month dollar interest rate, and the spot exchange rate. Also, we have already derived a formula for the

US dollar mark-to-market value of the forward by decomposing it into a long position in a pound denominated zero coupon bond with face value of £10 million and short position in a dollar denominated zero coupon bond with face value of \$15 million.

Step 2. The next step is to obtain historical values of the market factors for the last N periods. For our portfolio, this means collect the 3-month dollar and pound interbank interest rates and the spot dollar/pound exchange rate for the last 100 business days. Daily changes in these rates will be used to construct hypothetical values of the market factors used in the calculation of hypothetical profits and losses in Step 3 because the daily value at risk number is a measure of the portfolio loss caused by such changes over a one day holding period, 20 May 1996 to 21 May 1996..

Step 3. This is the key step. We subject the current portfolio to the changes in market rates and prices experienced on each of the most recent 100 business days, calculating the daily profits and losses that would occur if comparable daily changes in the market factors are experienced and the *current* portfolio is 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 upon, but *not* equal to, the historical values of the market factors over the past 100 days. Rather, we calculate daily historical percentage changes in the market factors, and then combine the historical percentage changes with the current (20 May 1996) market factors to compute 100 sets of hypothetical market factors.⁷ These hypothetical market factors are then used 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 20 May to obtain 100 hypothetical daily profits and losses.

Table 1 shows the calculation of the hypothetical profit/loss using the changes in the market factors from the first business day of 1996, which is day 1 of the 100 days preceding 20 May 1996. We start by using the 20 May 1996 values of the market factors to compute the mark-to-market value of the forward contract on 20 May, 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 12/29/95 to 1/2/96. The actual values on 12/29/95 and 1/2/96, and the percentage changes, are shown in lines 2 through 4. Then, in lines 5 and 6, we use the values of the market factors on 5/20/96, together with the percentage changes from 12/29/95 to 1/2/96, to compute hypothetical values of the market factors for 5/21/96. These hypothetical values of the market factors on 5/21/96 are then used to compute a mark-to-market value of the forward contract for 5/21/96 using the formula

⁷ This procedure of using the 20 May 1996 market factors together with the historical changes in order to generate hypothetical 21 May 1996 market factors makes sense because it guarantees that the hypothetical 21 May 1996 values will be more or less centered around the 20 May values, which is reasonable because the 20 May daily value at risk is a measure of the potential 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 will frequently involve using levels of the market factors that are not close to the current values. This reasoning, however, doesn't imply that one must use percentage changes together with the 20 May values in order to compute the hypothetical values of the market factors. 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 upon their levels.

$$\text{USD mark - to - market value} = \left[S \times \frac{\text{GBP 10 million}}{1 + r_{\text{GBP}}(90 / 360)} \right] - \frac{\text{USD 15 million}}{1 + r_{\text{USD}}(90 / 360)}.$$

This value is also shown on line 6. Once the hypothetical 5/21/96 mark-to-market value has been computed, the profit or loss on the forward contract is just the change in the mark-to-market value from 5/20/96 to 5/21/96, shown in line 7.

This calculation is repeated 99 more times, using the values of the market factors on 5/20/96 and the percentage changes in the market factors for days 2 through 100 to compute 100 hypothetical “mark-to-market” values of the forward contract for 5/21/96, and 100 hypothetical mark-to-market profits or losses. Table 2 shows these 100 daily mark-to-market profits and 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 are shown in Table 3, and range from a profit of \$212,050 to a loss of \$143,207.

Step 5. Finally, we select the loss which is equaled or exceeded 5 percent of the time. Since we have used 100 days, this is the fifth worst loss, or the loss of \$97,230, and is shown surrounded by a box on Table 3. Using a probability of 5 percent, this is the value at risk.

Figure 1 which was discussed previously shows the distribution of hypothetical profits and losses, with the value at risk indicated by an arrow. On the graph, the value at risk is the loss that leaves 5 percent of the probability in the left hand tail.

Multiple instrument portfolios

Extending the methodology to handle realistic, multiple instrument portfolios requires only that a bit of additional work be performed in three of the steps. First, in Step 1 there are likely to be many more market factors, namely the interest rates for longer maturity bonds and the interest and exchange rates for many 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 on each of the last N periods, or else by treating the volatilities as constants and disregarding the fact that they change randomly over time. This has the potential of introducing significant errors for portfolios with significant options content. Second, in Step 2 the historical values of all of the market factors must be collected. Third, it is crucial 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 value at risk is intended to capture the fact that typically gains on some instruments offset losses on others. Netting the gains against the losses within each of the 100 days in Step 3 reflects this relationship.⁸

⁸ The alternative procedure of ordering the profits and losses on the individual instruments before summing them to obtain the portfolio profits and losses implicitly assumes that the profits and losses on the individual instruments are perfectly positively correlated and usually results in a value at risk number that overstates the potential portfolio loss.

What determines the value at risk?

In order to understand the next methodology, it is useful to discuss the determinants of the value at risk in the simple example above. The value at risk of \$97,230 was determined by using the magnitudes of past changes in the market factors or their variability, the number of contracts in the portfolio (which was simply 1), the size of the forward contract (i.e., the quantities of dollars and pounds to be exchanged), and the sensitivity of its mark-to-market value to daily changes in the market factors. The number of forward contracts and its size translate into the face values of the zero coupon bonds into which it was decomposed, while the sensitivity of its value to changes in the market factors is captured by the sensitivities of the zero coupon bonds. The role of each of these is straightforward. More variable market factors, greater numbers of contracts, larger contracts, and contracts with greater sensitivities all result in a greater value at risk.

The value at risk is also determined by the comovement between the changes in the prices of the zero coupon bonds into which it was decomposed, or the extent to which changes in the value of the long position in the pound denominated bond are offset by changes in the value of the short position in the dollar denominated bond. This is determined by the extent to which dollar and pound interest rates, and the dollar/pound exchange rate, move together.

Variance-covariance approach

The variance/covariance approach is based on the assumption that the underlying market factors have a multivariate Normal distribution.⁹ Using this assumption (and other assumptions detailed below), it is possible to determine the distribution of mark-to-market portfolio profits and losses, which is also 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 value at risk.

For example, suppose we continue with our example of a portfolio consisting of a single instrument, the 3-month FX forward contract introduced above, and also continue to assume that the holding period is one day and the probability is 5%. The distribution of possible profits and losses on this simple portfolio can be represented by the probability density function shown in Figure 2. This distribution has a mean of zero, which is reasonable because the expected change in portfolio value over a short holding period is almost always close to zero. The standard deviation, which is a measure of the “spread” or dispersion of the distribution, is approximately \$52,500. A standard property of the Normal distribution is that outcomes less than or equal to 1.65 standard deviations below the mean occur only 5 percent of the time. That is, if a probability of 5 percent is used in determining the value at risk, then the value at risk is equal to 1.65 times the standard deviation of changes in portfolio value. Using this fact,

$$\begin{aligned}\text{value at risk} &= 1.65 \times \left(\begin{array}{l} \text{standard deviation of} \\ \text{change in portfolio value} \end{array} \right) \\ &= 1.65 \times 52,500 \\ &= 86,625.\end{aligned}$$

⁹ The name “variance-covariance” refers to the variance-covariance (or simply covariance) matrix of the distribution of changes in the values of the underlying market factors. An alternative name is the “analytic” method.

This value at risk is also shown in Figure 2. From this, it should be clear that the computation of the standard deviation of changes in portfolio value is the focus of the approach.

While the approach may seem rather like a “black box” because it is based on just a handful of formulas from statistics textbooks, it captures the determinants of value at risk mentioned above. It identifies the intuitive notions of variability and comovement with the statistical concepts of standard deviation (or variance) and correlation. These determine the variance-covariance matrix of the assumed Normal distribution of changes in the market factors. The number and size of the forward contract are captured through the “risk mapping” procedure discussed below. Finally, the sensitivity of the values of the bonds which comprise the instruments to changes in the market factors is captured in Step 4.

Risk mapping

A key step in the variance covariance approach is known as “risk mapping.” This involves taking the actual instruments and “mapping” them into a set of simpler, standardized positions or instruments. Each of these standardized positions is associated with a single market factor. For example, for the 3-month forward contract the basic market factors are the three month dollar and pound interest rates, and the spot exchange rate. The associated standardized positions are a dollar denominated 3-month zero coupon bond, a 3-month zero coupon bond exposed only to changes in the pound interest rate (i.e., it as if the exchange rate were fixed), and spot pounds. The covariance matrix of changes in the values of the standardized positions can be computed from the covariance matrix of changes in the basic market factors.¹⁰ This is illustrated in Step 3 below. Once the covariance matrix of the standardized positions has been determined, the standard deviation of any portfolio of the standardized positions can be computed using a single formula for the standard deviation of a sum of Normal random variables.¹¹

The difficulty is that the formula applies only to portfolios of the standardized positions. This creates the need for risk mapping. In order to compute the standard deviation and value at risk of any other portfolio, it must first be “mapped” into a portfolio of standardized positions. In essence, for any actual portfolio one finds a portfolio of the standardized positions that is (approximately) equivalent to the original portfolio in the sense that it has the same sensitivities to changes in the values of the market factors. One then computes the value at risk of that equivalent portfolio. If the set of standardized positions is reasonably rich and the actual portfolio doesn’t include too many options or option-like instruments then little is lost in the approximation.

Performing the analysis for a single instrument portfolio

We again illustrate the various steps involved using a portfolio consisting of a single instrument, the 3-month FX forward contract to deliver \$15 million on the delivery date 91 days hence, and in exchange receive £10 million. The method requires 4 steps.

¹⁰ The designer of the risk measurement system may choose the standardized positions to be the basic market factors, in which case this step isn’t necessary.

¹¹ The change in the value of a portfolio is the sum of the changes in the values of the positions which comprise it, so the standard deviation of changes in the value of a portfolio is the standard deviation of a sum.

Step 1. The first step is to identify the basic market factors and the standardized positions that are 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. We use a simple set of standardized positions in order to illustrate the procedure. A natural choice corresponds to our previous decomposition of the forward contract into a long position in a 3-month pound denominated zero coupon bond with a face value of £10 million and short position in a 3-month dollar denominated zero coupon bond with a face value of \$15 million. As indicated above, we take the standardized positions to be 3-month dollar-denominated zero coupon bonds, 3-month pound denominated zero coupon bonds that are 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 dollar leg and a pound leg, we have already completed a good bit of the work involved in mapping the contract. We need only to finish the process.

The dollar leg of the forward contract is easy. The value of a short position in a dollar denominated zero coupon bond with a face value of \$15 million can be obtained by discounting using the dollar interest rate. Letting X_1 denote the number of dollars invested in the first standardized position and using a negative sign to represent a short position, we have

$$X_1 = -\frac{\text{USD 15 million}}{1+r_{\text{USD}}(91/360)} = -\frac{\text{USD 15 million}}{1+.05469(91/360)} = \text{USD -14,795}.$$

The pound leg must be mapped into two standardized positions because its value depends on two market factors, the 3-month pound interest rate and the spot dollar/pound exchange rate. The magnitudes of the standardized positions are determined by separately considering how changes in each of the market factors affects the value of the pound leg, holding the other factor constant. The dollar value of the pound leg is

$$\begin{aligned} \text{dollar value of pound leg} &= (S \text{ USD / GBP}) \times \frac{\text{GBP 10 million}}{1+r_{\text{GBP}}(91/360)} \\ &= (15355 \text{ USD / GBP}) \times \frac{\text{GBP 10 million}}{1+.06063(91/360)} \\ &= \text{USD 15,123,242.} \end{aligned}$$

Holding the spot exchange rate S constant, this has the risk of $X_2 = 15,123,242$ dollars invested in 3-month pound bonds. Holding the pound interest rate constant, the bond with a face value of GBP 10 million has the exchange rate risk of a spot position of $\frac{\text{GBP 10 million}}{1+.06063(91/360)}$ pounds (its present value), or \$15,123,242. Hence the dollar value of the spot pound position is $X_3 = 15,123,242$. The equality of X_2 and X_3 is not coincidence, because both represent the dollar value of the pound leg of the forward contract. The dollar value of the pound leg of the contract appears twice in the mapped position because, from the perspective of a US company, a position in a pound denominated bond is exposed to changes in two market risk factors.

Having completed this mapping, the forward contract is now described by the magnitudes of the three standardized positions, X_1 , X_2 , and X_3 . Appendix B sketches a mathematical argument which justifies this mapping.

Step 2. The second step is to assume that percentage changes in the basic market factors have a multivariate Normal distribution with means of zero, and estimate the parameters of that distribution. This is the point at which the variance-covariance procedure captures the variability and comovement of the market factors: variability is captured by the standard deviations (or variances) of the Normal distribution, and the comovement by the correlation coefficients. The estimated standard deviations and correlation coefficients are shown in Table 4.

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 value 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% when the first market factor changes by 1%, then its standard deviation is twice as large as the standard deviation of the first 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 market factor. For example, the correlation between the first and third market factors, the dollar interest rate and the dollar/pound exchange rate, is 0.19, while 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 dollar interest rate. Appendix B formalizes this discussion.

Step 4. Now that we have the standard deviations of and correlations between changes in the values of the standardized positions, we can calculate the portfolio variance and standard deviation using standard mathematical results about the distributions of sums of Normal random variables and determine the distribution of portfolio profit or loss. The variance of changes in mark-to-market portfolio value depends upon the standard deviations of changes in the value of the standardized positions, the correlations, and the sizes of the positions, and is given by the standard formula

$$\begin{aligned}\sigma_{\text{portfolio}}^2 = & X_1^2 \sigma_1^2 + X_2^2 \sigma_2^2 + X_3^2 \sigma_3^2 + 2X_1 X_2 \rho_{12} \sigma_1 \sigma_2 \\ & + 2X_1 X_3 \rho_{13} \sigma_1 \sigma_3 + 2X_2 X_3 \rho_{23} \sigma_2 \sigma_3\end{aligned}$$

The standard deviation is of course simply the square root of the variance. For our example, the portfolio standard deviation is approximately $\sigma_{\text{portfolio}} = 52,500$.

One property of the Normal distribution is that outcomes less than or equal to 1.65 standard deviations below the mean occur only 5 percent of the time. That is, if a probability of 5 percent is used in determining the value at risk, then the value at risk is equal to 1.65 times the portfolio standard deviation. Using this, we can calculate the value at risk:

$$\begin{aligned}
\text{value at risk} &= 1.65 \times \sigma_{\text{portfolio}} \\
&= 1.65 \times 52,500 \\
&= 86,625.
\end{aligned}$$

As was discussed above, Figure 2 shows the probability density function for a Normal distribution with a mean of zero and a standard deviation of 52,500, along with the value at risk.

Realistic multiple instrument portfolios

Using a 3-month forward contract in the example allowed us to sidestep one minor difficulty. 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? It seems natural to write a formula for its value in terms of the 4-month U.S. dollar and British pound interest rates, just as we did with the 3-month forward. But doesn't this introduce two more market factors, the 4-month dollar and pound interest rates?

The answer is no. The 1, 3, 6, and 12 month interest rates are natural choices for market risk factors because there are active interbank deposit markets at these maturities, and rates for these maturities are widely quoted. In a number of currencies there are also liquid government bond markets at some of these maturities. There isn't an active 4-month interbank market in the U.S. dollar, the British pound, or any other currency. As a result, the 4-month interest rates used in computing the model value of the 4-month forward would typically be interpolated from the 3 and 6-month interest rates. (The interpolated 4-month rates might also depend on rates for the other actively quoted maturities, depending upon the interpolation scheme used.) Through this process, the current mark-to-market values of all dollar/pound forward contracts, regardless of delivery date, will depend on the spot exchange rate and the interest rates at only a limited number of maturities. As a result, value at risk measures computed using theoretical pricing models depend upon only a limited number of basic market factors.

The 4-month forward just mentioned could be handled as follows. We suppose that the forward price is 1.5 \$/£, and that the contract requires a U.S. company to deliver \$15 million and receive £10 million in four months. The first step is to decompose the forward contract into pound and dollar denominated 4-month zero coupon bonds just as we did with the 3-month forward. Next, the 4-month zeros must be "mapped" onto the 3 and 6-month zeros. The idea is to replace each of the 4-month zeros with a portfolio of the 3 and 6-month standardized positions that has the same market value and risk, where here "risk" means standard deviation of changes in mark-to-market value, which is proportional to value at risk. An instrument with multiple cash flows at different dates, for example a 10-year gilt, would be handled by mapping the 20 semi-annual cash flows onto the 6 and 12-month, and 2, 3, 4, 5, 7, 9, and 10-year pound denominated zero coupon bonds, the standardized positions. Each cash flow would be mapped onto the two nearest standardized positions.

The second section of Appendix C uses the 4-month dollar denominated zero to illustrate one way to perform this mapping. Appendix C also describes how options are mapped into their "delta-equivalent" standardized positions.

Relatively minor complications of realistic portfolios are that standard deviations and correlations must be estimated for all of the market factors, and the portfolio variance must be calculated using the appropriate generalization of the formula used above.

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, one chooses a statistical distribution that is believed to adequately capture or approximate the possible changes in the market factors. Then, a psuedo-random number generator is used to generate thousands or perhaps tens of thousands of hypothetical changes in the market factors. These are then 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 value at risk is then determined from this distribution.

A single instrument portfolio

Once again, we use the same portfolio of a single 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. This has already been done: the market factors are the 3-month pound interest rate, the 3-month dollar interest rate, and the spot exchange rate, and we have already derived a formula for the mark-to-market value of the forward by decomposing it into a portfolio of dollar and pound denominated 3-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 to 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 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 4.

The assumed distribution need not be the multivariate Normal, though the natural interpretations of its parameters (means, standard deviations, and correlations) and the ease with which these parameters can be estimated weigh in its favor. The designers of the risk management system are free to choose any distribution that they think reasonably describes possible future changes in the market factors. Beliefs about possible future changes in the market factors are typically based on observed past changes, so this amounts to saying that the designers of the risk management system are free to chose any distribution that they think approximates the distribution of past changes in the market factors.

Step 3. Once the distribution has been selected, the next step is to use a psuedo-random generator to generate N hypothetical values of changes in the market factors, where N is almost certainly greater than 1000 and perhaps greater than 10,000. These hypothetical market factors are then

used to calculate N hypothetical mark-to-market portfolio values. Then from each of the hypothetical portfolio values we subtract the actual mark-to-market portfolio value on 20 May 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 value at risk is the loss which is equaled or exceeded 5 percent of the time.

Multiple instrument portfolios

Just as with historical simulation, extending the methodology to handle realistic, multiple instrument portfolios requires only that a bit of additional work be performed in three of the steps. First, in Step 1 there are likely to be 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. Again, options may be handled either by treating the option volatilities as additional market factors that must be simulated, or else treating the volatilities as constants and disregarding the fact that they change randomly over time. Second, in Step 2 the joint distribution of possible changes in the values of all of 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 it is necessary that the mark-to-market profits and losses on every instrument be computed and then summed for each day, before they are ordered from highest profit to lowest loss in Step 4.

WHICH METHOD IS BEST?

With three methods from which to choose, the obvious question is: which method of calculating value at risk is best? Unfortunately, there is no easy answer. The methods differ in their ability to capture the risks of options and option-like instruments, ease of implementation, ease of explanation to senior management, flexibility in analyzing the effect of changes in the assumptions, and reliability of the results. The best choice will be determined by which dimensions the risk manager finds most important. Below we discuss how the three methods differ on these dimensions, and Table 5 summarizes the differences. We also discuss a closely related issue, the choice of the holding period t .

It may be that the best choice is not to use value at risk at all. Nonfinancial corporations might find that value at risk's focus on mark-to-market profit or loss over a holding period of t days doesn't match their perspective. Rather, they may be more interested in the distributions of quarterly cash flow over the next perhaps 20 quarters, and how these distributions are affected by transactions in financial instruments. This suggests a "cash flow at risk" measure, which we briefly discuss below when we describe alternatives to value at risk. Finally, as described below, companies with exposures to only a few different market factors may find simple sensitivity analyses to be adequate.

Ability to capture 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 variance-covariance method works well for instruments and portfolios with limited options content but is less able to capture the risks of options and option-like instruments than the two simulation methods. The limitation of the variance-covariance method is that it incorporates options by replacing them with or mapping them to their “delta-equivalent” spot positions (see Appendix B). This amounts to linearizing the options positions, or replacing the nonlinear functions which give their values in terms of the underlying rates and prices with linear approximations. For instruments or portfolios with a great deal of options content, the linear approximations may not adequately capture how the values of the options change with changes in the underlying rates and prices.

In the variance-covariance method, the problem of adequately capturing the risks of options and option-like instruments is least severe when the holding period is one day ($t=1$). Large changes in the underlying rates or prices are unlikely over such a short holding period, and the linear approximation in this method works well for small changes in the underlying rates and prices. As a result, the variance-covariance method works well even for positions with moderate options content provided the holding period is short. However, over longer holding periods, for example two weeks or one month, larger changes in underlying rates and prices are likely and value at risk estimates produced using the variance-covariance method cannot be relied upon for positions with moderate or significant options content.

The simulation methods work well regardless of the presence of options in the portfolio because they recompute the value of the portfolio for each “draw” of the basic market factors. In doing this, they estimate the “correct” distribution of portfolio value, though this statement must be qualified. The distribution of portfolio value generated by Monte Carlo simulation depends upon the assumed statistical distribution of the basic market factors and the estimates of its parameters, both of which can be “wrong” and therefore lead to errors in the calculated value at risk. 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 value at risk methodologies to incorporate the fact that option volatilities are random and option prices change with changes in volatilities. As indicated previously, the variance-covariance method also does not capture these features of options very well. In contrast, Monte Carlo simulation can incorporate, in principle, the facts that volatilities are random and option prices change with volatilities by extending the simulation to include a distribution of volatilities, though this typically is not done in actual implementation of this methodology. 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

The historical simulation method is easy to implement for portfolios restricted to currencies for which data on the past values of the basic market factors are available. It is conceptually simple, and can be implemented in a spreadsheet because pricing models for financial products are now available as spreadsheet add-in functions. The principal difficulty in implementing historical simulation is that it requires that the user possess a time series of the relevant market factors covering the last N days or other periods. This can pose a problem for multinational companies with operations and local currency borrowing in many countries, or with receivables and other

instruments in a wide range of currencies. While spot exchange rates are readily available for virtually all currencies, obtaining reliable daily market interest rates for a range of maturities in some currencies without well developed capital markets can be difficult.

A range of vendors offer software which computes value at risk estimates using the variance-covariance method, so this method is very easy to implement for portfolios restricted to currencies and types of instruments covered by the available systems. The variance-covariance method can be moderately difficult to implement for portfolios which include currencies and types of instruments not covered by the available systems. First, estimates of the standard deviations and correlations of the market factors are required. Computing these estimates is straightforward if data are available, but as indicated above reliable market interest rates may not be available for a range of maturities in all currencies.¹² Second, and more difficult, instruments must be mapped to the delta-equivalent positions as described in Appendix B.

“Off the shelf” software is starting to become available for the Monte Carlo simulation method, making it as easy to implement as the variance-covariance method for portfolios covered by the available systems. One difference is that computation times will be longer with Monte Carlo simulation. For portfolios not covered by the existing software, Monte Carlo simulation is in some ways easier, and in some ways more difficult, than the variance/covariance method. It is easier because it is not necessary to map instruments onto the standard positions, and it is more difficult because the user must select the distribution from which the psuedo-random vectors are drawn, and select or estimate the parameters of that distribution. Actually carrying out the simulation is not difficult because psuedo-random number generators are available as spreadsheet add-ins. However, selecting the distribution and selecting or estimating the parameters requires high degrees of expertise and judgment. Another disadvantage of Monte Carlo simulation is that it for large portfolios the computations can be time consuming.

All three methods require that pricing models be available for all instruments in the portfolio.¹³ While the variance/covariance 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 portfolios which included certain exotic options and currency swaps with complex embedded options.

Ease of communication with senior management

The conceptual simplicity of historical simulation makes it easiest to explain to senior management. The variance-covariance method is difficult to explain because to an audience without technical training because the key step, the reliance on the mathematics of the Normal distribution to calculate the portfolio standard deviation and the value at risk, is simply a black box. Monte Carlo simulation is even more difficult to explain. The key steps of choosing a statistical distribution to represent changes in the market factors and engaging in psuedo-random sampling from that distribution are simply alien to most people.

¹² Problem is slightly less severe than with historical simulation.

¹³ However, the pricing models need not be perfect because value at risk 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.

Reliability of the results

All methods rely on historical data. Historical simulation is unique, though, in that it relies so directly on historical data. A danger in this is that the price and rate changes over last 100 (or 200) days is that the last 100 (or 200) 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 value at risk computed using historical simulation would understate the risk in the portfolio. Alternatively, if by chance the U.S. dollar price of the Mexican peso rose steadily over the last 100 days and there were relatively few days on which the dollar price of a peso fell, value at risk computed using historical simulation would indicate that long positions in the Mexican peso involved 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 value at risk is underestimated, and for which it is overestimated. If value at risk is used to set risk or position limits, the traders can exploit their knowledge of the biases in the value at risk system and expose the company to more risk than the risk management committee intended.

Other methodologies use historical data to estimate the parameters of distributions (for example the variance-covariance methodology relies on historical data to estimate the standard deviations and correlations of a multivariate Normal distribution of changes in market factors for which the means are assumed to be zero), and are also subject to the problem that the historical period used might be atypical. However, assuming a particular distribution inherently limits the possible shapes that the estimated distribution can have. For example, if one assumed that the changes in the U.S. dollar price of a Mexican peso followed a Normal distribution with a mean of zero, one would predict that there was a 50 percent chance that the price of a peso would fall tomorrow even if the price had risen on each of the last 100 days. Since 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 is likely a better prediction than the prediction implicit in historical simulation.

The variance-covariance and Monte Carlo simulation methods share a different potential problem: the assumed distributions might not adequately describe the actual distributions of the market factors. Typically, actual distribution of changes in market rates and prices have “fat tails” relative to the Normal distribution. That is, there are more occurrences away from the mean than predicted by a Normal distribution. Nonetheless, the Normal distribution assumed in the variance-covariance method appears to be a reasonable approximation for the purposes of computing value at risk.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 of the system to make a bad choice, in the sense that the chosen distribution might not adequately approximate the actual distribution of the market factors.

Concerns about the reliability of the methods can be partially addressed by comparing actual changes in value to the value at risk amounts. This sort of validation is feasible because the value at risk approach explicitly specifies the probability with which actual losses will exceed the value at risk amount. It is performed by collecting a sample of value at risk amounts and actual market-to-market portfolio profits and losses, and answering two questions. First, does the distribution of

¹⁴ A good discussion of this issue may be found in J.P Morgan’s *RiskMetrics - Technical Document*.

actual mark-to-market profits and losses appear similar to the distribution used to determine the value at risk amount? And second, do the actual losses exceed the value at risk 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 and losses to differ somewhat from the expected distribution. Because of this, reliable inferences about the quality of the value at risk estimates can only be made using by comparing relatively large samples of value at risk 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 value at risk amounts, because it will take many years to collect a large sample of monthly or quarterly value at risk amounts and portfolio profits and 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 the future ones. For example, in the period immediately prior to the departure of the British pound from the European Monetary System (EMS) in September 1992, the historical correlation between changes in the dollar/pound and dollar/mark exchange rates was very high. Yet a risk manager might have suspected that the pound would leave the EMS, and therefore that the correlation would be much lower in the future. How easily could she have calculated the value at risk in this “what-if” scenario using each of the three methods?

Historical simulation is directly tied to the historical changes in the basic market factors. As a result, there is no natural way to perform this sort of “what-if” analysis.¹⁵ In contrast, it is very easy to carry out this sort of “what-if” analysis in the variance-covariance and Monte Carlo simulation methods. In these, the historical data are used to estimate the parameters of the statistical distribution of changes in the market factors. The user may override the historical estimates, and use any consistent set of parameters she chooses. The only constraint is that the user interfaces in some software implementations of the methods may make this cumbersome.

SUPPLEMENTING VALUE AT RISK: STRESS TESTING AND SCENARIO ANALYSIS

Value at risk is not a panacea. It is a single, summary, statistical measure of normal market risk. If a probability of 5 percent and a holding period of 1 day are used in computing the value at risk, you expect to suffer a loss exceeding the value at risk 1 (business) day out of 20, or about once per month. A level of loss that will be exceeded about once per month is reasonably termed a “normal” loss. But when the value at risk is exceeded, just how large can the losses be?

Stress testing attempts to answer this question. It is a general rubric for performing a set of scenario analyses to investigate the effects of extreme market conditions. To the extent that the effects are unacceptable, the portfolio or risk management strategy needs to be revised. There is

¹⁵ In this method, 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. Handling alternative assumptions about the correlations between a market factor and each of the others is possible, but considerably more cumbersome.

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 stylized 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 upon the changes in US dollar interest rates and bond prices experienced during the winter and spring of 1994, or the dramatic changes in some of the European exchange rates that occurred in September 1992. Alternatively, the scenarios might be created by imagining a few sudden surprises, and thinking through the implications for the markets. For example, how would the unanticipated failure of a major dealer affect prices and liquidity in the currency swaps market? What would be the effect on the Korean won and the Japanese yen if the North Koreans crossed the 38th parallel? What would be the effect of such an incident on the U.S. and Japanese equity markets? In developing these scenarios, it is important to think through the implications for all markets. An event sufficiently significant to have a sudden, major impact on the dollar/yen exchange rate would almost certainly impact other exchange rates, and likely affect interest rates in many currencies. A full description of a scenario will include the changes in all market rates and prices.

After developing a set of scenarios, the next step is to determine the effect on the prices of all instruments in the portfolio, and the impact on portfolio value. In addition, companies whose risk management strategies depend upon “dynamic hedging” or the ability to frequently adjust or rebalance their portfolios need to consider the impact of major surprises on market liquidity. It may be difficult or impossible to execute transactions at reasonable bid/ask spreads during periods of market stress. Companies which 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 contracts are received or paid immediately, while 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 combining futures contracts with other instruments can lead to timing mismatches between when funds are required and when they are received.

Finally, contingency plans might be developed for certain of the scenarios. Declines in market value, once suffered, typically cannot be recovered, so contingency plans have little to offer in this dimension. However, potential funding mismatches created by the cash demands of futures positions can be managed by arranging backup lines of credit. The potential importance of this is illustrated by MG Refining and Marketing (MGRM), a classic example of a firm which was not prepared to meet the funding demands of its futures positions. MGRM is a U.S. subsidiary of Metallgesellschaft A.G., the 14th largest German industrial firm, and was engaged in the refining and marketing of petroleum products in the United States. Among its activities, MGRM used futures contracts and short-term commodity swaps on crude oil and various refined products to hedge long-term delivery obligations. In early 1994 it had to be rescued by a group of 150 German and international banks when it was unable to meet the funding needs created by staggering losses on its futures contracts and swaps. Regardless of one’s view on the wisdom of using futures to hedge long-term delivery obligations and MGRM’s risk management strategy,¹⁶

¹⁶ The wisdom of MGRM’s hedging strategy and the parties primarily to blame for the losses have been the subject of considerable controversy. Views generally supportive of MGRM’s risk management strategy and critical of the parent management’s response to the difficulties are expressed by Christopher L. Culp and Merton H. Miller in a number of papers: “Metallgesellschaft and the Economics of Synthetic Storage,” *Journal of Applied Corporate Finance* 7, No. 4 (Winter 1995), pp. 62-76; “Hedging a flow of Commodity Derivatives with Futures: Lessons from

in retrospect it seems clear that MGRM's failures included the lack of a plan for meeting the funding demands of its futures contracts.

Scenario analyses are also used to examine the effects of violations of the assumptions underlying the value at risk calculations. For example, immediately prior to the British pound's departure from the EMS in September 1992, all three value at risk methodologies would have indicated that from the perspective of a U.S. dollar investor a long position in sterling combined with a short position in Deutsche marks had a very low value at risk. The low value at risk would have been a result of the historically high correlations between the dollar/pound and dollar/mark exchange rates, for all three value at risk methodologies rely upon historical data. Yet in September 1992 the position would have suffered a large loss, because the historical correlations could no longer be relied upon. This risk could be evaluated either by changing the correlation used as an input in calculating the value at risk, or by examining directly the impact on portfolio if the pound fell relative to the mark. Regardless, the key input to this process is the risk manager's judgment that the scenario is worth considering.

ALTERNATIVES TO VALUE AT RISK

As indicated above, value at risk may not be appropriate for all entities. Two alternatives are sensitivity analysis and cash flow at risk. Sensitivity analysis is less sophisticated than value at risk. In contrast, cash flow at risk can be considered more sophisticated than value at risk.

Sensitivity analysis

Companies with exposures to only a few market factors may find that the benefits of value at risk don't justify the difficulty of mastering the approach and implementing a system to compute the value at risk estimates. As discussed next, sensitivity analyses are a reasonable alternative for sufficiently simple portfolios.

The approach in sensitivity analysis is to imagine hypothetical changes in the value of each market factor, and then use pricing models to compute the value of the portfolio given the new value of the market factor and determine the change in portfolio value resulting from the change in the market factor. For example, if the dollar price of a pound increases by 1%, the value of the portfolio will decrease by \$200,000; if the dollar price of a pound decreases by 1%, the value of the portfolio will increase by \$240,000. There is nothing magical about 1%. Rather, the computations will typically be performed and reported for a range of increases and decreases that cover the range of likely exchange rate changes. Similar computations would also 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 very good picture of the risks of portfolios with

Metallgesellschaft," *Derivatives Quarterly* 1, No. 1 (Fall 1994), pp. 7-15; "Auditing the Auditors," *Risk* 8, No. 4 (April 1995); and "Hedging in the Theory of Corporate Finance: A Reply to Our Critics," *Journal of Applied Corporate Finance* 8, No. 1 (Spring 1995), pp. 121-8. Contrary views are expressed by Antonio Mello and John Parsons in "Maturity Structure of a Hedge Matters: Lessons from the Metallgesellschaft Debacle" *Journal of Applied Corporate Finance* 8, No. 1 (Spring 1995), pp. 106-20 and by Franklin R. Edwards and Michael S. Canter in "The Collapse of Metallgesellschaft: Unhedgeable Risks, Poor Hedging Strategy, or Just Bad Luck," *Journal of Applied Corporate Finance* 8, No. 1 (Spring 1995), pp. 86-105.

exposures to only a few market factors. In fact, they comprise the most basic risk management information, and are very closely related to the delta risk measure discussed in Appendix A. In one form or another, market risk sensitivities have been available to traders and risk managers since at least 1938.¹⁷ 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, each representing the change in portfolio value for a particular hypothetical change in market rates and prices. Absent some approach like value at risk, it is difficult or impossible for a risk manager or senior manager charged with oversight of trading and risk management activities to meaningfully read and review sensitivity analysis reports for portfolios with exposures to many different market factors and assimilate the information to get a sense of portfolio risk.¹⁸

Cash flow at risk

As stated previously, cash flow at risk is arguably more sophisticated than value at risk. As of this writing, it appears to have a limited, but growing, number of users. Cash flow at risk is a reasonable choice for nonfinancial corporations which are concerned with managing the risks inherent in operating cash flows and find that value at risk's focus on mark-to-market profit or loss over a holding period of t days doesn't match their perspective.

For example, Merck is a user of both derivatives and cash flow at risk. The motivation for derivatives usage appears to be the fact that changes in cash flows due to changes in interest and exchange rates were negatively impacting R&D programs by causing shortfalls of funds.¹⁹ Currency and interest rate swaps, appropriately used, are able to ameliorate this problem. But this motivation for derivatives usage suggests that the risk measurement system ought to focus on quarterly or annual cash flows over a horizon of at least several years. For example, a company in a similar situation might be interested in the distributions of quarterly cash flow over the next perhaps 20 quarters, and how these distributions are affected by transactions in financial instruments.²⁰

Cash flow at risk measures are typically estimated using Monte Carlo simulation. However, there are important differences from the use of Monte Carlo simulation to estimate value at risk. First, the time horizon is much longer in cash flow at risk simulations. For example, values of

¹⁷ The concept of the duration of a bond was invented by Frederick Macaulay in 1938 (*Some Theoretical Problems Suggested by the Movements of Interest Rates, Bond Yields, and Stock Prices in the United States Since 1865*, National Bureau of Economic Research). Macaulay duration is closely related to modified duration, which is a sensitivity expressed in percentage terms.

¹⁸ Alternatively, a portfolio sensitivity analysis calculation could be performed assuming that all market risk factors change by given percentages simultaneously. However, this joint sensitivity to multiple changes in market factors also suffers in comparison to value at risk because it does not ensure that equally likely losses are aggregated across different classes of instruments. While of course one can add the profit or loss stemming from an x percent change in dollar interest rates to the profit or loss stemming from a y percent change in the dollar/yen exchange result, it isn't clear that the resulting sum has any meaningful interpretation.

¹⁹ See Judy C. Lewent and A. John Kearny, Identifying, Measuring, and Hedging Currency Risk at Merck, in Donald H. Chew, Jr., *The New Corporate Finance: Where Theory Meets Practice*, McGraw-Hill, 1993, 452-461.

²⁰ In contrast, a company whose Treasury group actively manages a portfolio of borrowing, swaps, and other interest and exchange rate instruments, perhaps in order to exploit perceived profit opportunities or trends in market rates and prices, would be more likely to find value at risk useful. Some corporations might use both methods.

the underlying market factors might be simulated for the next 20 quarters. Second, the focus is on cash flows, not changes in mark-to-market values. This is the distinguishing feature, and in fact the whole point, of cash flow at risk measures. Rather than using the hypothetical values of the market factors as inputs to pricing models to compute changes in mark-to-market portfolio value, the hypothetical market factors are combined with the terms of the cash and derivative instruments to compute hypothetical quarterly or annual cash flows, and their distributions. Third, *operating* cash flows are typically included in the calculation. This is of course essential if the goal of the risk measurement system is to assess the impact of derivatives and other financial transactions on companies' total cash flows. As a result, the "factors" included in the simulation are not just the basic financial market factors included in value at risk calculations, but any "factors" which affect operating cash flows. Changes in customer demand, the outcomes of R&D programs (including competitors' R&D programs), and competitors' pricing decisions are a few operating "factors" that come to mind. Finally, the emphasis is often on planning rather than control, oversight, and reporting.

A serious drawback is that successful design and implementation of a cash flow at risk measurement system requires a high degree of knowledge and judgment.²¹ First, the designer of the system must develop a model of the company's operating cash flows, determining the important operating factors and how they impact operating cash flows. This alone may be a major undertaking. Next, this model of the operating cash flows must be integrated with a model of the 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 can be particularly difficult for the operating "factors." In contrast with 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 the difficulty of the problem is that the model of the financial market factors can be relatively crude, as there is no point in refining it to be more precise than the model of the operating cash flows. Nonetheless, building a cash flow at risk measurement system is likely to be a major undertaking.

²¹ "Off-the-shelf" software is currently not available, and may never be available because cash flow at risk systems typically include operating cash flows, the characteristics of which are company-specific and difficult to incorporate in an "off-the-shelf" system. However, at least one major derivatives dealer has been willing to provide some current and potential future customers with the framework of a cash flow at risk system, the simulation engines, and assistance in implementing the system.

APPENDIX A

BASIC RISK MEASURES: OPTION DELTA'S AND GAMMA'S

Delta

The delta or Δ is perhaps the most basic risk management concept. Delta indicates how much the theoretical price of an instrument or portfolio changes when the price of the underlying asset, currency, or commodity changes by a small amount. Therefore it is very closely related to sensitivity analysis. While originally developed for options, the concept can be applied to other derivatives, and to cash positions as well.

We illustrate the concept of delta using a call option on British pounds with a strike price of 1.50 \$/£ and 3 months to expiration. We suppose that the current dollar/pound exchange rate is also 1.50 \$/£ and the current price of the call option is \$0.0295 per pound. The price of this option will vary as the dollar/pound exchange rate varies. Figure 4 shows the theoretical price (computed using the Garman-Kohlhagen model) as a function of the dollar/pound exchange rate. The graph indicates that if the dollar/pound exchange rate changes slightly from the current value of 1.50 \$/£, the change in the option price will be about one-half as large as the change in the exchange rate. For example, if the exchange rate changes to 1.51 \$/£, the (theoretical) option price will change by \$0.0051 to \$0.0346. The ratio of the change in the option price to the

change in the currency price, $\frac{\$0.0051}{\$0.01} = 0.51$, is the option delta. Graphically, the delta is the slope of the line which is tangent to the option price function at the current exchange rate. This tangent is shown in Figure 4. Formally, delta is the partial derivative of the option price function with respect to the underlying currency price. Letting S denote the dollar price of a British pound and $C(S)$ denote the option price as a function of S , the option delta is

$$\Delta \equiv \frac{\partial C(S)}{\partial S}.$$

Since delta is given by the ratio of price changes, i.e.

$$\Delta = \frac{\text{change in option price}}{\text{change in price of underlying instrument}},$$

the change in the option price resulting from a change in the spot price can be calculated from the delta and the change in the price of the underlying instrument:

$$\text{change in option price} = \Delta \times \text{change in price of underlying instrument}$$

For example, if $\Delta = 0.51$ and the price of a pound changes by \$0.01, the predicted change in the option price is $\$0.0051 = 0.51 \times \0.01 . One interpretation of this relationship is that an option on one pound is equivalent to a spot position of delta British pounds, because the change in value of a spot position of delta British pounds is also given by the product of delta and the change in the spot price of a pound. Loosely, for small changes in the exchange rate the option “acts like” delta British pounds. The significance of this for risk measurement is that one technique for

measuring the risk of an option position is to use the option delta to compute the equivalent spot position, and then estimate the risk of the equivalent spot position. Most applications of the variance-covariance methodology for computing value at risk which we discuss below rely on this technique.

An important feature of options and option-like instruments is that delta changes as the price of the underlying asset, currency, or commodity changes. This is illustrated in Figure 5, which shows the

theoretical price of a 3-month call option on pounds with a strike price of 1.50 \$/£, together with the option deltas. At the current spot price of 1.50 \$/£ the delta is approximately one-half, while for high spot prices the delta approaches one and for low spot prices it approaches zero. The delta approaches one for high spot prices because if the spot price is well above the strike price the option is almost certain to be exercised. An option that is almost certain to be exercised behaves like a levered position in the underlying asset or currency. The delta approaches zero for low spot prices because if the spot price is well below the strike price the option is almost certain to expire unexercised. An option that is almost certain to expire unexercised is worth almost nothing now, and behaves like almost nothing.

The changing delta illustrated in Figure 5 doesn't appear to be a severe problem for risk measurement. However, for many options positions reliance solely on delta can be misleading. Figure 6 shows the value of one such position as a function of the dollar/pound exchange. The portfolio shown in Figure 6 consists of a spot position in 1 pound along with 2 written 3-month options. At the spot exchange rate 1.50 \$/£, the delta of the spot pound is 1 and the delta of the call option is approximately 0.5, so the portfolio delta is approximately $1 - 2 \times 0.5 = 0$.²² Using a delta of zero to compute the equivalent spot position, we would conclude that this options position is equivalent to a spot position of zero British pounds, and therefore has no market risk. But clearly the position does have market risk, for if the exchange rate changes in either direction by more than a small amount the position will suffer a loss.

Gamma

Gamma or Γ supplements delta by measuring how delta changes as the price of the underlying asset, currency, or commodity changes. In Figure 6 delta decreases as the dollar price of a pound increases, so gamma is negative. (The slope is positive for \$/£ exchange rates less than 1.50 \$/£, and negative for exchange rates greater than 1.52 \$/£.) If delta increases as the dollar price of a pound increases, then gamma is positive. 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 $C(S)$ denote the option price as a function of S , the option gamma is

²² The delta of a cash position in the underlying asset, currency, or commodity is always 1, because when the "derivative" and the underlying instrument are identical the ratio of the change in the price of the derivative to the change in the price of the underlying instrument is simply 1.

$$\Gamma \equiv \frac{\partial \Delta(S)}{\partial S}$$

$$= \frac{\partial^2 C(S)}{\partial S^2}.$$

Delta and gamma together can be used to predict the change in the option price resulting from a change in the spot price of one pound using the following formula:

$$\text{change in option price} = \Delta \times \left(\text{change in price of} \right) + \frac{1}{2} \Gamma \times \left(\text{change in price of} \right)^2$$

Comparing this to the earlier equation which predicts the change in the option price using only delta, one can see that when gamma is negative the change in the option price is more adverse than that predicted using delta alone. Conversely, when gamma is positive the change in the option price is more favorable than that predicted using delta alone

The significance of this for value at risk measures is that the variance-covariance method typically measures the risk of options by converting them to their equivalent spot positions using delta alone and thereby somewhat understate the risk of positions with negative gammas. The effect will be small for value at risk computations done using short holding periods, because for short holdings periods the change in the spot price of the underlying asset is typically small and the term

$$\frac{1}{2} \Gamma \left(\text{change in price of} \right)^2$$

is small. However, the understatement of the risk of negative gamma portfolios can be significant when value at risk measures are computed for long holding periods.

APPENDIX B

CALCULATION OF STANDARD DEVIATIONS AND CORRELATIONS OF PERCENTAGE CHANGES IN THE VALUES OF THE STANDARDIZED POSITIONS

In essence, if the value of the standardized position changes by x percent when the market factor changes by 1 percent, then the standard deviation of percentage changes in the standardized position is equal to x times the standard deviation of percentage changes in the market factor. To see this more formally, let X_1 denote the value of the first standardized position, and use the fact that

$$\begin{aligned}\% \text{ change in } X_1 &\approx \frac{\partial X_1}{\partial r_{\text{USD}}} \times \frac{1}{X_1} \times \text{change in } r_{\text{USD}} \\ &= \frac{\partial X_1}{\partial r_{\text{USD}}} \times \frac{r_{\text{USD}}}{X_1} \times \% \text{ change in } r_{\text{USD}}\end{aligned}$$

This implies that

$$\text{std. deviation of \% change in } X_1 \approx -\frac{\partial X_1}{\partial r_{\text{USD}}} \times \frac{r_{\text{USD}}}{X_1} \times \text{std. deviation of \% change in } r_{\text{USD}},$$

where the minus sign appears because $\frac{\partial X_1}{\partial r_{\text{USD}}}$ is negative, i.e., the value of the first standardized position moves inversely with USD interest rates. Letting σ_1 denote the standard deviation of percentage changes in X_1 and σ_{USD} denote the standard deviation of percentage changes in the dollar interest rate, this can be rewritten

$$\sigma_1 \approx -\frac{\partial X_1}{\partial r_{\text{USD}}} \times \frac{r_{\text{USD}}}{X_1} \times \sigma_{\text{USD}}.$$

Similarly, for the other two standardized positions:

$$\begin{aligned}\sigma_2 &\approx -\frac{\partial X_2}{\partial r_{\text{GBP}}} \times \frac{r_{\text{GBP}}}{X_2} \times \sigma_{\text{GBP}}, \\ \sigma_3 &\approx \frac{\partial X_3}{\partial S} \times \frac{S}{X_3} \times \sigma_S.\end{aligned}$$

In addition, the signs of two of the correlation coefficients must be changed because the values of the first and second standardized positions move inversely with the USD and GBP interest rates. Due to this, we have $\rho_{13} = -\rho_{\text{USD},S}$, and $\rho_{23} = -\rho_{\text{GBP},S}$. The correlation between the first two standardized positions is unaffected because both move inversely with interest rates, and $\rho_{12} = \rho_{\text{USD},\text{GBP}}$.

APPENDIX C

RISK MAPPING

Theory underlying mapping the forward contract into the three standardized positions

Here we show that the forward contract can be described as a portfolio of the three standardized positions with the same sensitivities to the market factors. In other words, they have the same risks. This is the key to risk mapping. We do this 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 choose the standardized positions so that the coefficients of the two Taylor series approximations are the same. If the coefficients of the Taylor series 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

$$V_F \equiv S \times \left[\frac{\text{GBP 10 million}}{1 + r_{\text{GBP}}(91/360)} \right] - \frac{\text{USD 15 million}}{1 + r_{\text{USD}}(91/360)}$$

denote the mark-to-market value of the forward contract. Using a Taylor series, the change in V_F can be approximated

$$\Delta V_F \approx \frac{\partial V_F}{\partial r_{\text{USD}}} \Delta r_{\text{USD}} + \frac{\partial V_F}{\partial r_{\text{GBP}}} \Delta r_{\text{GBP}} + \frac{\partial V_F}{\partial S} \Delta S.$$

Next, we will write down a similar Taylor series approximation of changes in the value of the portfolio of standardized positions, and show that if the standardized positions are chosen appropriately then the coefficients of the two approximations are identical. If this is true then $\Delta V \approx \Delta V_F$, implying 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

$$\Delta V \approx \frac{\partial X_1}{\partial r_{\text{USD}}} \Delta r_{\text{USD}} + \frac{\partial X_2}{\partial r_{\text{GBP}}} \Delta r_{\text{GBP}} + \frac{\partial X_3}{\partial S} \Delta S.$$

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 amounts to choosing them so that

$$\frac{\partial X_1}{\partial r_{\text{USD}}} = \frac{\partial V_F}{\partial r_{\text{USD}}}, \quad \frac{\partial X_2}{\partial r_{\text{GBP}}} = \frac{\partial V_F}{\partial r_{\text{GBP}}}, \quad \text{and} \quad \frac{\partial X_3}{\partial S} = \frac{\partial V_F}{\partial S}.$$
 The choice that works is

$$X_1 = -\frac{\text{USD 15 million}}{1 + r_{\text{USD}}(91 / 360)},$$

$$X_2 = \frac{(1.5355 \text{ USD / GBP}) \times \text{GBP 15 million}}{1 + r_{\text{GBP}}(91 / 360)},$$

$$X_3 = (S \text{ USD / GBP}) \frac{\text{GBP 15 million}}{1 + 0.06063(91 / 360)}.$$

These are three standardized positions we used before to carry out the risk mapping of the forward contract. As indicated earlier, they are interpreted as follows. The first, X_1 , is simply the value of a position in 3-month dollar denominated bonds. The other two are more complicated. X_2 is the dollar value of the position in 3-month pound denominated bonds, holding the exchange rate fixed, while X_3 is the dollar value of a spot position in pounds equal to the present value of the pound bonds, holding the pound interest rate fixed. Note that both X_2 and X_3 represent the value of the pound denominated bond, but each of them is exposed to only one of the two market factors that affect the value of the bond.

Mapping a 4-month dollar denominated cash flow onto the 3 and 6-month standardized positions

The idea is to replace the 4-month cash flow with a portfolio of the 3 and 6-month standardized positions that has the same risk or distribution of changes in market value as the original cash flow. This requires that the portfolio has the same market value and standard deviation (or variance) of changes in market value.

To find the market value of the original 4-month cash flow, we need an interest rate with which to discount it. One way to obtain a 4-month US dollar interest rate is simply to interpolate using the 3 and 6-month rates. This amounts to taking the 4-month rate to be a weighted average of the 3 and 6-month rates, or

$$r_{4\text{-mo}} = (2 / 3)r_{3\text{-mo}} + (1 / 3)r_{6\text{-mo}}.$$

The present value of the dollar leg of the 4-month forward is then

$$PV = -\frac{\text{USD 15 million}}{1 + (1 / 3)r_{4\text{-mo}}},$$

where the 1/3 appears in the denominator because the cash flow must be discounted for one-third of a year.

The standard deviation of changes in the value of the 4-month cash flow depends upon the sensitivity of changes its value to changes in the interest rate and the standard deviation of changes in the interest rate. In symbols,

$$\sigma_{PV} = \frac{\partial PV}{\partial r_{4\text{-mo}}} r_{4\text{-mo}} \sigma_{4\text{-mo}},$$

where $\frac{\partial PV}{\partial r_{4\text{-mo}}}$ is the sensitivity of changes in the value of the dollar leg to changes in the interest rate, $\sigma_{4\text{-mo}}$ is the standard deviation of percentage changes in the 4-month rate, and $r_{4\text{-mo}} \sigma_{4\text{-mo}}$ is the standard deviation of (“absolute”) changes in the 4-month rate. The parameter $\sigma_{4\text{-mo}}$ can be computed from the 3 and 6-month rates, the standard deviations of percentage changes in the 3 and 6-month rates, and the correlation between these changes using standard results for linear combinations of Normal random variables.

Next, introduce a fourth standardized position consisting of 6-month dollar denominated zero coupon bonds, and let X_4 denote the value of the position. The mapping of the 4-month cash flow onto the 3 and 6-month standardized positions is completed by finding a portfolio of X_1 dollars in 3-month bonds and X_4 dollars in 6-month bonds. This portfolio must have the same value and standard deviation of changes in value as the 4-month cash flow. Also, the signs of X_1 and X_4 must be the same as the sign of the 4-month cash flow. In symbols, we need to find a portfolio X_1 and X_4 such that:

$$\begin{aligned} PV &= X_1 + X_4, & (\text{values match}) \\ \sigma_{PV} &= \text{standard deviation} (X_1 + X_4), & (\text{standard deviations match}) \\ \text{sign}(X_1) &= \text{sign}(X_4) = \text{sign}(-15 \text{ million}). & (\text{signs match}) \end{aligned}$$

The last equation is needed because the first two equations will typically have two different solutions for X_1 and X_4 , one of which will involve a negative sign. The standard deviation of the portfolio with value $X_1 + X_4$ is computed using the technique discussed in Step 3 of the section on the variance-covariance method. Finally, these equations are solved for X_1 and X_4 .

Mapping Options

Options positions typically are mapped into “delta equivalent” positions in spot foreign currency and the standardized zero coupon bonds. An option delta is the partial derivative of the option price with respect to the price of the underlying asset. Letting V denote the theoretical value of the option and S denote the price of the underlying asset, the delta is

$$\Delta = \frac{\partial V}{\partial S}.$$

As discussed more fully in Appendix A, the change in the option price resulting from a change in the spot price can be calculated from the delta and the change in the price of the underlying asset:

$$\text{change in } V = \Delta \times \text{change in } S.$$

For example, if the option is on 1 million British pounds, $\Delta = 0.5$ million or 0.5 per pound, and the spot price of one pound changes by \$0.01, the predicted change in the option price is $\$0.005 = 0.5 \times \0.01 million. One interpretation of the equation above is that for small changes in the exchange rate an option is equivalent to a spot position of Δ British pounds, because the change in value of a spot position of Δ British pounds is also given by the product of Δ and the change in the spot price of 1 pound. Loosely, the option “acts like” Δ British pounds.

Mapping of other options positions is conceptually the same, though sometimes more complicated. Consider an over-the-counter option on a 10-year British gilt. Usually, one would say that the underlying asset is a 10-year gilt. However, recall that we indicated that the 20 semi-annual cash flows of a 10-year gilt might be mapped onto the 6 and 12-month, and 2, 3, 4, 5, 7, 9, and 10-year pound denominated zero coupon bonds. If we took the perspective of a pound investor, we would interpret the option on the gilt as an option on a portfolio of these 9 zero coupon bonds, and think of the option as having 9 underlying assets and nine deltas, one for each underlying asset. However, the dollar price of the gilt also depends on the dollar/pound exchange rate. From the perspective of a dollar investor, there are 10 underlying assets: the nine pound denominated zero coupon bonds, along with the dollar/pound exchange rate, and for each we can define a delta. Letting V denote the dollar value of the gilt and P_n denote the pound price of the n th pound denominated zero, for the first nine deltas we have

$$\Delta_n = \frac{\partial V}{\partial P_n} \quad .$$

The tenth delta, the partial derivative with respect to the spot exchange rate, is

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

The change in the option price resulting from changes in the prices of the underlying assets is given by

$$\text{change in } V = \sum_{n=1}^9 (\Delta_n \times \text{change in } P_n) + \Delta_{10} \times \text{change in } S \quad .$$

The change in V is identical to the change in the value of a portfolio of Δ_n units of each of the nine pound zeros, along with Δ_{10} spot pounds. Exploiting this observation, the option is “mapped” into this portfolio.

To understand why this procedure can be useful, remember that value at risk is a portfolio level risk measure. It is computed by assigning a risk measure to each position, and then aggregating up to a portfolio level measure. A difficulty is that there are an immense variety of different options. Even if we just consider ordinary options, wide ranges of both strike prices and expiration dates are possible, and of course there are both calls and puts. In addition, there are exotic options which can have virtually any terms. How can one reasonably assign a risk measure to every option? The approach in most variance-covariance value at risk systems is to measure the risk of a set of standardized positions, and then measure the risk of options in terms of the delta-equivalent positions.

Explicit risk mapping of this sort is only necessary in the “analytic” or “variance-covariance” methodology. However, in this framework it is the key issue in the design of a value at risk system. To hint at the complexities, consider a second option, but this time suppose it is a futures option on the British pound currency futures traded on the International Money Market (IMM) of the Chicago Mercantile Exchange. It seemed natural to map the first option on spot pounds into a Δ - equivalent spot position. Should the IMM futures option also be mapped into a Δ - equivalent spot position by using the theoretical relationship between currency spot and futures prices to reinterpret it as an option on spot pounds? Or should we introduce a second basic market risk factor, the futures price, and map the futures option into a Δ - equivalent futures position? What if we consider another futures option on a pound futures contract with a different delivery date? And what about the fact that option and futures prices change with changes in interest rates? The answers to these questions are not obvious. Nonetheless, the questions need to be answered by the designer of a value at risk system.

Table 1: Calculation of Hypothetical 5/21/96 Mark-to-Market Profit/Loss on a Forward Contract Using Market Factors from 5/20/96 and Changes in Market Factors from the First Business Day of 1996

	Market Factors			Mark-to-Market Value
	\$ Interest Rate (% per year)	£ Interest Rate (% per year)	Exchange Rate (\$/£)	of Forward Contract (\$)
Start with actual values of market factors and forward contract as of close of business on 5/20/96:				
(1) Actual values on 5/20/96	5.469	6.063	1.536	327,771
Compute actual past changes in market factors:				
(2) Actual values on 12/29/95	5.688	6.500	1.553	
(3) Actual values on 1/2/96	5.688	6.563	1.557	
(4) Percentage change from 12/29/95 to 1/2/96	0.000	0.962	0.243	
Use these to compute hypothetical future values of the market factors and the mark-to-market value of the forward contract:				
(5) Actual values on 5/20/96	5.469	6.063	1.536	327,771
(6) Hypothetical future values calculated using rates from 5/20/96 and percentage changes from 12/29/95 to 1/2/96	5.469	6.121	1.539	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

$$\text{USD mark - to - market value} = \left[(\text{exchange rate in USD / GBP}) \times \frac{\text{GBP 10 million}}{1 + r_{\text{GBP}}(90 / 360)} \right] - \frac{\text{USD 15 million}}{1 + r_{\text{USD}}(90 / 360)}.$$

Table 2: Historical Simulation of 100 Hypothetical Daily Mark-to-Market Profits and Losses on a Forward Contract

Number	Market Factors			Hypothetical Mark-to-Market Value of Forward Contract (\$)	Change in Mark- to-Market Value of Forward Contract (\$)
	\$ Interest Rate (% per year)	£ Interest Rate (% per year)	Exchange Rate (\$/£)		
1	5.469	6.121	1.539	362,713	34,942
2	5.379	6.063	1.531	278,216	-49,555
3	5.469	6.005	1.529	270,141	-57,630
4	5.469	6.063	1.542	392,571	64,800
5	5.469	6.063	1.534	312,796	-14,975
6	5.469	6.063	1.532	294,836	-32,935
7	5.469	6.063	1.534	309,795	-17,976
8	5.469	6.063	1.534	311,056	-16,715
9	5.469	6.063	1.541	379,357	51,586
10	5.438	6.063	1.533	297,755	-30,016
.					
.					
.					
91	5.469	6.063	1.541	378,442	50,671
92	5.469	6.063	1.545	425,982	98,211
93	5.469	6.063	1.535	327,439	-332
94	5.500	6.063	1.536	331,727	3,956
95	5.469	6.063	1.528	249,295	-78,476
96	5.438	6.063	1.536	332,140	4,369
97	5.438	6.063	1.534	310,766	-17,005
98	5.469	6.125	1.536	325,914	-1,857
99	5.469	6.001	1.536	338,368	10,597
100	5.469	6.063	1.557	539,821	212,050

Table 3: Historical Simulation of 100 Hypothetical Daily Mark-to-Market Profits and Losses on a Forward Contract, Ordered From Largest Profit to Largest Loss

Number	Market Factors			Hypothetical Mark-to-Market	Change in Mark-to-Market Value of Forward Contract (\$)
	\$ Interest Rate (% per year)	£ Interest Rate (% per year)	Exchange Rate (\$/£)	Value of Forward Contract (\$)	
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

Table 4: Standard Deviations of and Correlations Between % Changes in Market Factors

Market Factor	Standard Deviations of % Changes	Correlations Between % Changes in Market Factors			
		Market Factor	3-month \$ interest rate	3-month £ interest rate	\$/£ exchange rate
3-month \$ interest rate	0.61	3-month \$ interest rate	1.00		
3-month £ interest rate	0.58	3-month £ interest rate	0.11	1.00	
\$/£ exchange rate	0.35	\$/£ exchange rate	0.19	0.10	1.00

Table 5: Comparison of Value at Risk Methodologies

	Historical Simulation	Variance/Covariance	Monte Carlo Simulation
Able to capture the risks of portfolios which include options?	Yes, regardless of the options content of the portfolio	No, except when computed using a short holding period for portfolios with limited or moderate options content	Yes, regardless of the options content of the portfolio
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 upon 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.
Computations performed quickly?	Yes.	Yes.	No, except for relatively small portfolios.
Easy to explain to senior management?	Yes.	No.	No.
Produces misleading value at risk 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.
Easy to perform “what-if” analyses to examine effect of alternative assumptions?	No.	Easily able to examine alternative assumptions about correlations/standard deviations. Unable to examine alternative assumptions about the distribution of the market factors, i.e. distributions other than the Normal.	Yes.

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




Figure 2: Probability Density Function and Value at Risk Obtained Using Variance-Covariance Method




Figure 3: Focus of “Stress Testing”




Figure 4: Price and Delta of a Call Option on British Pounds




Figure 5: Delta Changes as the Exchange Rate Changes




Figure 6: Example of a Risky Portfolio that has Delta = 0

