The quest for precision through Value at Risk

by Paul Glasserman

Discussions of financial risk can bring to mind the story of the six blind men describing an elephant. One man feels the trunk and concludes the elephant is like a snake. Another touches a leg and likens the beast to a tree. Each of the blind men manages to state something strictly correct, yet none of their descriptions fully captures all aspects of the creature.

Financial risk is at least as difficult to summarize as an elephant. The difficulty is not one of complete ignorance: like the blind men in the story, we can grope our way at least to some partially correct conclusions. Yet we are rarely able to give a complete picture of one of the most fundamental drivers of economic behavior.

At the risk of unfair oversimplification, the evolution of financial risk management over the past five years might be summarized as a quest to describe risk succinctly through a single number, Value at Risk or VaR. The quest has been championed by regulatory bodies, aided by a string of well-publicised financial disasters, fed by a growing software and consulting industry, and widely embraced by the financial companies and corporations exposed to the risk everyone is trying to measure. To what extent has the quest been successful? What lessons have been learned along the way?

We need to start by defining Value at Risk or VaR, which, briefly, is a percentile of a profit-and-loss distribution over a specified horizon. To give some life to that definition, let us consider an example. Suppose we have responsibility for measuring the risk in a portfolio of financial assets; the portfolio could consist of currencies, bonds, equities, derivatives, or a combination of these. We fix a time horizon over which to consider the risk – two weeks, for example. We might ask what the most is we could possibly lose on this portfolio over this period. The answer would be the worst-case loss. By definition, we can be 100 percent certain that losses from the portfolio will not exceed that amount. However, the worst-case loss gives too pessimistic a picture of risk, because the worst possible outcome is unlikely to occur. Instead, we might therefore ask for an amount (in dollars or pounds, say) such that we can be 95 percent certain that losses will not exceed that amount. The answer to that question is the Value at Risk. More precisely, it is the VaR for a two-week horizon and a 95 percent confidence level.

In reporting a VaR number, we understand that it is not a maximum loss. Indeed, if we regularly measure VaR at a 95 percent confidence level, then actual losses should exceed our estimate 5 percent of the time. If that is uncomfortable, we can use a 99 percent VaR and thus only underestimate 1 percent of the time. The point is that by not insisting on a 100 percent confidence level, we are providing a more realistic picture of losses a portfolio might suffer.

To appreciate the strengths and weaknesses of VaR, we should contrast it with other measures of risk. A simple but crude measure of the risk in a financial position is the outright exposure. A \$10m position in government bonds has an exposure of \$10m. So does a \$10m position in internet stocks. While it is theoretically possible to lose $10m in either case, large losses are certainly more likely in the second one. VaR tries to capture this.

Compare a $10m position in three-month government bonds with a $10m position in 30-year government bonds. The two are exposed to the same source of risk (US interest rates), but fixed-income managers have long understood that the longer maturity carries greater risk, because the value of long-term bonds is more sensitive to changes in interest rates. Sensitivity to interest-rate changes is thus an important risk measure in fixed-income markets. Sensitivities are also commonly used to measure and hedge the risks in options.

A difficulty in using sensitivities to measure risk is that they cannot be compared across markets and products. Duration measures, for example, are useful in comparing the risks in bonds issued by a single government, but they are much less useful in making international comparisons or comparing risks in bonds with other asset classes. In contrast, a VaR number is always calculated as an amount. A VaR of $55m is intended to measure the same level of risk, regardless of whether it applies to a portfolio of bonds, currencies, stocks, or options.

Students of modern portfolio theory are familiar with the statistical concept of standard deviation as a measure of risk. Standard deviation is used to measure the size of statistical fluctuations of many phenomena and the application to changes in portfolio value is natural. Like VaR, the standard deviation blends information about the probabilities of losses together with the size of losses. The most important difference between these two risk measures is that standard deviation implicitly assumes that profits and losses are mirror images of each other: a $1m loss has the same likelihood as a $1m gain, and the same for any other dollar amount. While this may be approximately true for simple instruments – say, a cash position in a major foreign currency – it is far from true when applied to options. Buying an option creates unlimited upside potential with no chance of losing more than the cost of the option. Conversely, writing options can produce unlimited downside risk with no possibility of profit beyond the premium earned on the sale of the option. By focussing only on the chance of large losses, VaR can capture this sort of asymmetry whereas standard deviation cannot.

A brief history

Two events have had a particularly strong impact on the widespread adoption of VaR in the financial sector and a third has encouraged its growth in the US corporate sector.

The first took place in Basle, Switzerland, in 1995. Meeting at the Bank for International Settlements, a committee of representatives from the central banks of 10 major western economies proposed new rules (amending a 1988 Basle accord) requiring financial institutions to hold capital against their exposure to market risk. The proposal, formally adopted in 1996, created an incentive for banks to develop sophisticated internal risk systems to calculate their VaR. Banks doing so could expect to see a reduction in the risk capital they would be required to hold to sustain trading activities, compared with banks that followed a standardized approach imposed by regulators to determine capital requirements. Thus, from the outset, regulatory relief has been an important factor in the growth of VaR.

The second important event took place on the World Wide Web. Starting in 1994, the US bank J.P. Morgan made its RiskMetrics system available to all over the internet at no cost. RiskMetrics (since then spun off in a joint venture with Reuters, the financial information group) provided financial data and a methodology to calculate a portfolio's Value at Risk. Other financial institutions and corporations were free to use the RiskMetrics VaR Calculator or download the RiskMetrics data to their own risk management systems. Third-party vendors quickly emerged providing risk management software tapping into RiskMetrics, making the methodology an instant benchmark.

The third event has probably had less total impact to date, but is one of the leading factors driving the use of VaR among US corporations. In 1997, the US Securities and Exchange Commission, concerned about undisclosed risks lurking in off-balance-sheet instruments, issued disclosure rules for the use of derivative securities by corporations. The rules allow corporations to choose one of three possible ways of disclosing risks from derivatives use: a tabulation of fair market value, a sensitivity measure, or VaR. Consequently, one can now read about VaR calculations in the annual reports of Microsoft, consumer products manufacturer Philip Morris, and many other major corporations.

Market realities

The approach to VaR popularized by J.P. Morgan is a good example of how it can be more effective to move aggressively with an imperfect solution than to wait until all the wrinkles have been smoothed out before taking action. The RiskMetrics methodology is based on a series of assumptions that simplify the calculation of VaR but do not always hold in practice. The developers of the methodology understood this; by now, most of its users probably do too. However, this simple methodology has usefully framed the discussion over how VaR should be calculated, and has helped focus efforts both in industry and academia on essential obstacles to estimating VaR accurately. Many of the important difficulties in determining VaR are best understood by contrasting market realities with these simplifying assumptions.

Markets are not normal

One of the central assumptions of the RiskMetrics approach and of many other VaR models is that market returns can be described by the statistical notion of a normal distribution. This is the familiar bell-shaped curve used to describe phenomena ranging from the heights of British prisoners to scores on school exams. A graphical record of returns on actively traded market instruments does at first glance bear resemblance to a normal curve: most returns fall in the middle near the average; very large moves up or down occur less frequently.

However, on closer inspection, market data deviates from the mathematical ideal of a normal distribution in a remarkably consistent way. Across virtually all major markets and asset classes, returns exhibit what statisticians (and increasingly risk managers) call "excess kurtosis." In simple terms, this means that too many observations occur near the center of the distribution – the curve is too peaked. But these values near the average are inevitably accompanied by an excess of very extreme values.

Imagine taking a bell-shaped lump of clay and squeezing the middle. Some of the clay gets pushed up, creating a higher peak, but some of it gets pushed down and out to the sides. This part corresponds to extreme market moves. Compared with the ideal of a normal distribution, the risk manager sees a world in which market prices move either too little or too much. This effect is particularly pronounced in measuring returns over short time horizons (for example, a day or a week). It is so pervasive that it comes close to qualifying as a law of nature in the financial world.

Portfolios are non-linear

Many VaR calculations assume that portfolio value moves in strict proportion to changes in market prices. This is often true. If you buy 100 shares of a stock at $23, then if the stock price increases by $1 your holdings increase in value by $100. If the stock price increases by $2 you make $200. The relation between the value of the portfolio and the market price of the stock can be described by a straight line. However, this type of linear relation often breaks down for portfolios that include derivatives. For example, the value of a call option on the stock might increase by $0.48 when the stock moves from $23 to $24 and then by $0.52 when the stock moves from $24 to $25. The relation between the stock price and the value of the option can no longer be described by a straight line. For portfolios consisting primarily of cash instruments, this may have little impact. However, for portfolios that include substantial positions in options, failing to capture this effect can render a VaR estimate meaningless.

The author's research, in collaboration with Philip Heidelberger of International Business Machines and Perwez Shahabuddin of Columbia University, has focussed on methods for estimating VaR that accurately capture the non-linear impact of options as well as the excess kurtosis characteristic of market data. This leads to more reliable VaR estimates, but it also puts greater data and computing demands on risk management systems.

Volatility is not constant

This may hardly seem surprising. Even a casual observer of financial news will note that markets appear more volatile at some times than at others. In risk management, volatility is used to describe a precise measure of market fluctuations (closely related to the standard deviation discussed above). While we know that volatility is not constant, accurately describing how volatility changes remains a challenge. Market data shows signs of "burstiness" or intermittency in volatility, with periods of high volatility engendering further volatility. A risk manager, therefore, cannot rely on a historical average level of volatility, but must forecast volatility based on current market conditions.

Volatility is also the key determinant of option value, so the difficulties in measuring volatility become particularly pronounced in calculating VaR for portfolios with options. The celebrated Black–Scholes formula, first published by economists Fischer Black and Myron Scholes in 1973, remains the cornerstone of option valuation. Nevertheless, it has been widely observed that market prices of options – especially on major equity indices such as the S&P 500 in the US – systematically deviate from this formula precisely because the formula assumes a constant level of volatility. Options that pay off in the event of a sharp decline in the S&P 500 are overpriced (according to a strict application of the Black–Scholes formula) relative to options that pay off in less extreme scenarios. Of course, when model and market disagree, the market is right and the model is wrong. Risk managers have become highly sensitive to this limitation of option theory; they have to be in order to see the risk in what would otherwise look like a profitable strategy of selling crash insurance on the S&P 500.

Markets move together, but nobody knows how

Many of the greatest financial risks arise from co-movements among market risk factors. In the summer of 1997, several Asian currencies experienced sharp drops in quick succession. The effect of the Russian default was felt in Latin American bond markets. An increase in the price of oil can adversely affect major stock markets. The most basic way to try to quantify the propensity of market rates to move together (or apart) is through the statistical concept of correlation. Accurate measurement of market correlations is indeed central to quantitative risk management; virtually all methods for calculating VaR rely at least implicitly on estimated correlations. However, it is now widely recognized that statistical correlation remains at best an imperfect description of the way markets move together when they make big moves. No measure of correlation could have captured the spread of the Asian currency crisis from the Thai bhat to the Korean won, for example. Sound risk management must therefore supplement VaR estimates with stress-testing scenarios outside the domain of statistical models.

It may be easier to meet the needs of regulators than the needs of the company

Risk management at many financial institutions (and to a lesser extent at non-financial corporations) entails two parallel sets of activities: providing the VaR reports required by regulators, and providing broader input on risk to the management of the company. These are never quite the same. Satisfying regulators, though perhaps onerous, is more straightforward, in part because their requirements apply to the company as a whole. Internally, a company needs to know not just how much risk capital is required for the entire company, but how this capital should be allocated among its units. This in turn raises questions concerning internal measurement of risk-adjusted performance. Does VaR provide an appropriate risk measure for determining risk-adjusted performance? Are profits earned on risky positions evidence of financial acumen or simply of risk taking? These types of questions remain an open challenge for the future of risk management.

Companies throughout the world have made enormous investments in the systems and staff needed to monitor and report VaR. Many have achieved the objective of complying with legal requirements or even obtaining regulatory relief through these investments. Has calculating VaR produced sounder risk management? Few companies would attest that this single number has by itself produced better business practices. However, the effort to measure VaR has undoubtedly brought a heightened awareness of fundamental issues in market risk management and of the data, systems, and expertise needed to monitor risk. The lessons learned along the way and the lessons still to be learned may well be more valuable than any VaR report. The blind men may never agree on a single description of the elephant, but they can surely learn much by trying.

Summary

Despite advances in finance and risk management, a satisfactory method for measuring the total financial risk faced by a business or bank at any time remains elusive, says **Paul Glasserman**. Value at Risk (VaR) is one attempt, and it has certainly helped people grapple with basic issues in market risk management. Regulatory concern over the use of derivatives has helped popularize this, as have the opportunities for risk capital savings by banks that develop sophisticated VaR measurement systems. After explaining the fundamental principles behind VaR and the landmarks in its history, the author goes on to outline some of the difficulties in its measurement. Questionable assumptions include a normal

market, the linearity of portfolio values, and constancy of volatility. VaR's usefulness, he concludes, lies more in the processes and systems it encourages than in the substance of the single number it provides.

Copyright © Paul Glasserman 2001

Suggested further reading

Alexander, C. (ed.) (1998) *Risk Management and Analysis*, Chichester: Wiley.
Duffie, D. and Pan, J. (1997) "An overview of Value at Risk," *Journal of Derivatives*, Spring, 7–49.
Glasserman, P., Heidelberger, P., and Shahabuddin, P. (2000) "Portfolio Value at Risk with heavy-tailed risk factors," PaineWebber Research Series, The Columbia University Graduate Business.
Jorion, P. (1997) *Value at Risk*, New York: McGraw-Hill.
Longerstaey, J. and Spencer, M. (1996) *RiskMetrics Technical Document*, www.riskmetrics.com/research/techdocs

Insurance and finance: new vehicles for driving value

by Neil A. Doherty

Risk management is about insurance and hedging. If a company is exposed to volatile cash flows and there is a set of costs associated with volatility, then an obvious way to control those costs is to reduce volatility, that is, to hedge the risk. The financial risk management literature has developed to reflect the two prongs of this proposition. On the one hand, researchers have asked why risk is costly to the company; simultaneously, they have sought to analyze and price existing hedging instruments and to derive new or derivative instruments to hedge new and exotic sources of risk.

However, hedging is not the only way a company can offset the cost of risk. If one understands the structural features of the company that cause risk to be a problem, then value can be created by adapting the structure of the company so that it is more robust to risk. For example, one reason risk is costly is that volatility increases the chance that any given company will become bankrupt. This, in turn, triggers a set of bankruptcy costs. A company is bankrupt when it is unable to meet its debt obligations. So the problem can be addressed by reducing the volatility (which reduces the probability of falling below a fixed debt obligation), or by keeping the risk and switching from debt to equity financing. Hedging and capital structure choices address the same corporate problem. As we address other reasons for risk being costly, we will see that the cost can be reduced by either reducing the risk or making the operation more resilient to a given level of risk.

A simple valuation model of the company

The starting point for identifying how risk management can create value is a simple valuation model of the company. The value of equity, V(E), is the value of the