YieldCurve.com

An introduction to Value-at-Risk

Learning Curve

September 2003

Value-at-Risk

The introduction of Value-at-Risk (VaR) as an accepted methodology for quantifying market risk is part of the evolution of risk management. The application of VaR has been extended from its initial use in securities houses to commercial banks and corporates, and from market risk to credit risk, following its introduction in October 1994 when JP Morgan launched RiskMetrics™. VaR is a measure of the worst expected loss that a firm may suffer over a period of time that has been specified by the user, under normal market conditions and a specified level of confidence. This measure may be obtained in a number of ways, using a statistical model or by computer simulation.

> *VaR is a measure of market risk. It is the maximum loss which can occur with* X% *confidence over a holding period of* n *days.*

VaR is the expected loss of a portfolio over a specified time period for a set level of probability. For example if a daily VaR is stated as £100,000 to a 95% level of confidence, this means that during the day there is a only a 5% chance that the loss the next day will be *greater* than £100,000. VaR measures the potential loss in market value of a portfolio using estimated volatility and correlation. The "correlation" referred to is the correlation that exists between the market prices of different instruments in a bank's portfolio. VaR is calculated within a given confidence interval, typically 95% or 99%; it seeks to measure the possible losses from a position or portfolio under "normal" circumstances. The definition of normality is critical and is essentially a statistical concept that varies by firm and by risk management system. Put simply however, the most commonly used VaR models assume that the prices of assets in the financial markets follow a normal distribution. To implement VaR, all of a firm's positions data must be gathered into one centralised database. Once this is complete the overall risk has to be calculated by aggregating the risks from individual instruments across the entire portfolio. The potential move in each instrument (that is, each risk factor) has to be inferred from past daily price movements over a given observation period. For regulatory purposes this period is at least one year. Hence the data on which VaR estimates are based should capture all relevant daily market moves over the previous year.

VaR is only a measure of a bank's risk exposure; it a tool for measuring market risk exposure. There is no one VaR number for a single portfolio, because different methodologies used for calculating VaR produce different results. The VaR number captures only those risks that can be measured in quantitative terms; it does not capture risk exposures such as operational risk, liquidity risk, regulatory risk or sovereign risk.

Assumption of normality
A distribution is described as *normal* if there is a high probability that any observation form the population sample will have a value that is close to the mean, and a low probability of having a value that is far from the mean. The normal distribution curve is used by many VaR models, which assume that asset returns follow a normal pattern. A VaR model uses the normal curve to estimate the losses that an institution may suffer over a given time period. Normal distribution tables show the probability of a particular observation moving a certain distance from the mean.

If we look along a normal distribution table we see that at -1.645 standard deviations, the probability is 5%; this means that there is a 5% probability that an observation will be at least 1.645 standard deviations below the mean. This level is used in many VaR models. (We will present an introduction to standard deviation and the normal distribution in a later *Learning Curve*).

Calculation methods
There are three different methods for calculating VaR. They are:

- the variance/covariance (or *correlation* or *parametric* method);
- historical simulation;
- Monte Carlo simulation.

Variance-covariance method
This method assumes the returns on risk factors are normally distributed, the correlations between risk factors are constant and the delta (or price sensitivity to changes in a risk factor) of each portfolio constituent is constant. Using the correlation method, the volatility of each risk factor is extracted from the historical observation period. Historical data on investment returns is therefore required. The potential effect of each component of the portfolio on the overall portfolio value is then worked out from the component's delta (with respect to a particular risk factor) and that risk factor's volatility.

There are different methods of calculating the relevant risk factor volatilities and correlations. Two alternatives are:

- simple *historic volatility*: this is the most straightforward method but the effects of a large one-off market move can significantly distort volatilities over the required forecasting period. For example if using 30-day historic volatility, a market shock will stay in the volatility figure for 30 days until it drops out of the sample range and correspondingly causes a sharp drop in (historic) volatility 30 days *after* the event. This is because each past observation is equally weighted in the volatility calculation;
- to weight past observations unequally: this is done to give more weight to recent observations so that large jumps in volatility are not caused by events that occurred some time ago. One method is to use exponentially-weighted moving averages.

Historical simulation method

The historic simulation method for calculating VaR is the simplest and avoids some of the pitfalls of the correlation method. Specifically the three main assumptions behind correlation (normally distributed returns, constant correlations, constant deltas) are not needed in this case. For historical simulation the model calculates potential losses using actual historical returns in the risk factors and so captures the non-normal distribution of risk factor returns. This means rare events and crashes can be included in the results. As the risk factor returns used for revaluing the portfolio are actual past movements, the correlations in the calculation are also actual past correlations. They capture the dynamic nature of correlation as well as scenarios when the usual correlation relationships break down.

Monte Carlo simulation method

The third method, Monte Carlo simulation is more flexible than the previous two. As with historical simulation, Monte Carlo simulation allows the risk manager to use actual historical distributions for risk factor returns rather than having to assume normal returns. A large number of randomly generated simulations are run forward in time using volatility and correlation estimates chosen by the risk manager. Each simulation will be different but in total the simulations will aggregate to the chosen statistical parameters (that is, historical distributions and volatility and correlation estimates). This method is more realistic than the previous two models and therefore is more likely to estimate VaR more accurately. However its implementation requires powerful computers and there is also a trade-off in that the time required to perform calculations is longer.

Correlation

Measures of correlation between variables are important to fund managers who are interested in reducing their risk exposure through diversifying their portfolio. Correlation is a measure of the degree to which a value of one variable is related to the value of another. The correlation coefficient is a single number that compares the strengths and directions of the movements in two instruments values. The sign of the coefficient determines the relative directions that the instruments move in, while its value determines the strength of the relative movements. The value of the coefficient ranges from -1 to +1, depending on the nature of the relationship. So if, for exmple, the value of the correlation is 0.5, this means that one instrument moves in the same direction by half of the amount that the other instrument moves. A value of zero means that the instruments are uncorrelated, and their movements are independent of each other.

Correlation is a key element of many VaR models, including parametric models. It is particularly important in the measurement of the variance (hence volatility) of a portfolio. If we take the simplest example, a portfolio containing just two assets, equation (1) below gives the volatility of the portfolio based on the volatility of each instrument in the portfolio (x and y) and their correlation with one another.

$$V_{port} = \sqrt{x^2 + y^2 + 2xy \cdot \rho(xy)} \qquad (1)$$

where

x is the volatility of asset x
y is the volatility of asset y
ρ is the correlation between assets x and y.

The correlation coefficient between two assets uses the covariance between the assets in its calculation. The standard formula for covariance is shown at (2):

$$Cov = \frac{\sum_{i=1}^{n}(x_i - \bar{x})(y_i - \bar{y})}{(n-1)} \qquad (2)$$

where the sum of the distance of each value x and y from the mean is divided by the number of observations minus one. The covariance calculation enables us to calculate the correlation coefficient, shown as (3):

$$r = Cov\frac{(1,2)}{s_1 ´ s_2} \qquad (3)$$

where

s is the standard deviation of each asset.

Equation (1) may be modified to cover more than two instruments. In practice correlations are usually estimated on the basis of past historical observations. This is an important consideration in the construction and analysis of a portfolio, as the associated risks will depend to an extent on the correlation between its constituents.

It should be apparent that from a portfolio perspective a positive correlation increases risk. If the returns on two or more instruments in a portfolio are positively correlated, strong movements in either direction are likely to occur at the same time. The overall distribution of returns will be wider and flatter, as there will be higher joint probabilities associated with extreme values (both gains and losses). A negative correlation indicates that the assets are likely top move in opposite directions, thus reducing risk.

It has been argued that in extreme situations, such as market crashes or large-scale market corrections, correlations cease to have any relevance, because all assets will be moving in the same direction. However under most market scenarios using correlations to reduce the risk of a portfolio is considered satisfactory practice, and the VaR number for diversified portfolio will be lower than that for an undiversified portfolio.

Simple VaR calculation

To calculate the VaR for a single asset, we would calculate the standard deviation of its returns, using either its historical volatility or *implied volatility*. If a 95% confidence level is required, meaning we wish to have 5% of the observations in the left-hand tail of the normal distribution, this means that the observations in that area are 1.645 standard deviations away from the mean. Consider the following statistical data for a government bond, calculated using one year's historical observations.

Nominal:	£10 million
Price:	£100
Average return:	7.35%
Standard deviation:	1.99%

The VaR at the 95% confidence level is 1.645 x 0.0199 or 0.032736. The portfolio has a market value of £10 million, so the VaR of the portfolio is 0.032736 x 10,000,000 or £327,360. So this figure is the maximum loss the portfolio may sustain over one year for 95% of the time.

We may extend this analysis to a two-stock portfolio. In a two-asset portfolio, we stated at (1) that there is a relationship that enables us to calculate the volatility of a two-asset portfolio; this expression is used to calculate the VaR, and is shown at (4):

$$Var_{port} = \sqrt{w_1^2 s_1^2 + w_2^2 s_2^2 + 2w_1 w_2 s_1 s_2 r_{1,2}} \qquad (4)$$

where

w_1 is the weighting of the first asset

w_2 is the weighting of the second asset

s_1 is the standard deviation or *volatility* of the first asset

s_2 is the standard deviation or volatility of the second asset

$r_{1,2}$ is the correlation coefficient between the two assets.

In a two-asset portfolio the undiversified VaR is the weighted average of the individual standard deviations; the diversified VaR, which takes into account the correlation between the assets, is the square root of the variance of the portfolio. In practice banks will calculate both diversified and undiversified VaR. The diversified VaR measure is used to set trading limits, while the larger undiversified VaR measure is used to gauge an idea of the bank's risk exposure in the event of a significant correction or market crash. This is because in a crash situation, liquidity dries up as market participants all attempt to sell off their assets. This means that the correlation relationship between assets cease to have any impact on a book, as all assets move in the same direction. Under this scenario then, it is more logical to use an undiversified VaR measure.

Although the description given here is very simple, nevertheless it explains what is the essence of the VaR measure; VaR is essentially the calculation of the standard deviation of a portfolio, which is the used as an indicator of the volatility of that portfolio. A portfolio exhibiting high volatility will have a high VaR number. An observer may then conclude that the portfolio has a high probability of making losses. Risk managers and traders may use the VaR measure to help them to allocate capital to more efficient sectors of the bank, as return on capital can now be measured in terms of return on risk capital. Regulators may use the VaR number as a guide to the capital adequacy levels that they feel the bank requires.

Further illustration of variance-covariance VaR

Consider the following hypothetical portfolio, invested in two assets, as shown in figure 1. The standard deviation of each asset has been calculated on historical observation of asset returns. Note that *returns* are returns of asset prices, rather than the prices themselves; they are calculated from the actual prices by taking the ratio of closing prices. The returns are then calculated as the logarithm of the price relatives. The mean and standard deviation of the returns are then calculated using standard statistical formulae. This would then give the standard deviation of daily price relatives, which is converted to an annual figure by multiplying it by the square root of the number of days in a year, usually taken to be 250.

Assets	**Bond 1**	**Bond 2**	
Standard deviation	11.83%	17.65%	
Portfolio weighting	60%	40%	
Correlation coefficient			0.647
Portfolio value			£10,000,000
Variance			0.016506998
Standard deviation			12.848%
95% c.i. standard deviations			1.644853
Value-at-Risk			0.211349136
Value-at-Risk £			£2,113,491

Figure 1: Two-asset portfolio VaR.

The standard equation (shown as (4)) is used to calculate the variance of the portfolio, using the individual asset standard deviations and the asset weightings; the VaR of the book is the square root of the variance. Multiplying this figure by the current value of the portfolio gives us the portfolio VaR, which is £2,113,491.

The RiskMetrics VaR methodology uses matrices to obtain the same results that we have shown here. This is because once a portfolio starts to contain many assets, the method we described above becomes unwieldy. Matrices allow us to calculate VaR for a portfolio containing many hundreds of assets, which would require assessment of the volatility of each asset and correlations of each asset to all the others in the portfolio. We can demonstrate how the parametric methodology uses variance and correlation matrices to calculate the variance, and hence standard deviation, of a portfolio. The matrices for the example in Figure 1 are shown in Choudhry (2001).

The variance-covariance method captures the diversification benefits of a multi-product portfolio because the correlation coefficient matrix used in the calculation. For instance if the two bonds in our hypothetical portfolio had a negative correlation the VaR number produced would be lower. To apply it, a bank would require data on volatility and correlation for the assets in its portfolio. This data is actually available from the RiskMetrics website (and other sources), so a bank does not necessarily need its own data. It may wish to use its own datasets however, should it have them, to tailor the application to its own use. The advantages of the variance-covariance methodology are that:

- it is simple to apply, and fairly straightforward to explain;
- datasets for its use are immediately available.

The drawbacks of the variance-covariance are that it assumes stable correlations and measures only linear risk; it also places excessive reliance on the normal distribution, and returns in the market are widely believed to have "fatter tails" than a true to normal distribution. This phenomenon is known as *leptokutosis*, that is, the non-normal distribution of outcomes. Another disadvantage is that the process requires *mapping*. To construct a weighting portfolio for the RiskMetrics tool, cash flows from financial instruments are mapped into precise maturity points, known as grid points. We will review this later in the chapter, however in most cases assets do not fit into neat grid points, and complex instruments cannot be broken down accurately into cash flows. The mapping process makes assumptions that frequently do not hold in practice.

Nevertheless the variance-covariance method is still popular in the market, and is frequently the first VaR method installed at a bank.

*   *   *

Reference

Choudhry, M., *The Bond and Money Markets: Strategy, Trading, Analysis*, Butterworth-Heinemann 2001