4.3.1c Appendix C: Value at Risk

Unit: Unit 1 – Treasury Strategy

Module: Module 4 – Financial risk

Date: 1 September 2015

Summary: Value at Risk is a mathematically derived measure of risk, based essentially on the volatility of and correlations between the asset classes in a portfolio. It has powerful computability, bringing risk together in one measure, but does have weaknesses

Keywords: maximum loss, holding period, probability, confidence level, implied volatility, portfolio, diversification, stress testing, scenario analysis

Contents

1 Introduction 1
2 Value at Risk 1
  2.1 Working with implied volatilities 4
  2.2 Applicability 4
3 Portfolio effect 6
  3.1 Concept 6
  3.2 Advantages and caveats 8

© Association of Corporate Treasurers

Learning outcomes:

1. Value at Risk (VaR) has a short history and stems from the financial sector's struggle with risk management
2. VaR and its approach forms part of the language and culture of risk management
3. VaR is the maximum expected loss over a given holding period with a given probability
4. The volatility assumed in calculations can be historic (from observation) or implied (from market pricing)
5. VaR has many shortcomings
6. VaR is often supplemented by stress testing
7. VaR can be run using Monte Carlo simulation
8. VaR is difficult to apply to many business risks

© Association of Corporate Treasurers

1 Introduction

A history of Value at Risk [see website link below] [1] is a really useful introduction to the concepts and problems of VaR. Although it was written in 2002, it highlights concerns raised for it as a measure of risk, well before the crisis of 2008 focused attention on the weaknesses of mathematical representations of reality. It highlights the rapid growth in risk in the financial sector and the struggles to deal with it.

While VaR does have these problems, it remains a useful member of the risk evaluation toolkit. Still the mechanisms of its approach are used in pitch books by bankers and pension advisers when advancing products aimed at reducing risk. Still we talk in terms of 1 in 200 years failure, or 95% confidence of maximum loss, all language from VaR.

This reading covers the basic mathematics of VaR. Treasurers should be fluent with its concepts to properly assess the statements made by bankers, risk managers and consultants.

2 Value at Risk

Baldly, value at risk (VaR) is a statistical technique which summarises for a financial risk:

- the maximum expected loss (or worsening)
- over a given holding period
- with a given probability

For instance, suppose a company has identified a dollar-sterling exposure worth GBP10m at current levels. *Assuming that* over time daily changes in dollar-sterling exchange rates are normally distributed, then if we calculate the standard deviation using daily data, we can estimate the maximum expected loss over a twenty-four hour period which will be exceeded, say, only 5% of the time, i.e., will not be exceeded 95% of the time. The VaR "answer" is a money amount – GBP25,000 in this case, say. What it means is that if the corporate reviews its exposures on a daily basis, then 95% of the time it won't lose more than GBP25,000 on its exposure between reviews – but 5% of the time it will (statistically). At its most basic this allows management to quantify the risk in this single holding (rightly or wrongly – see disadvantages below – but at least it generates a reproducible and auditable money amount rather than relying on subjective comfort levels). The VaR number is then the answer to "How much risk am I living with for the period until I review my exposure again?" Obviously, the technique can be extended to examine more sophisticated cases.

This is an example of a risk measurement where one measure of risk is produced and the response to the risk can be based around the number, by either taking action to reduce the number or to allow it to continue in the form of a risk "budget".

The ***holding period*** in this example is one day. The length of the holding period is usually based on how long it would take to extinguish the exposure, allowing for operational

[1] http://www-stat.wharton.upenn.edu/~steele/Courses/434/434Context/RiskManagement/VaRHistlory.pdf

mechanics. For instance, on a bank trading floor, VaR might be set on an hourly basis during the day, with an overnight overlay. For an illiquid fixed-rate bond the holding period would be longer and the risk of a price movement would be that much bigger. In the corporate context with (say) fortnightly or monthly reviews of exposures, the holding period would match the review period.

The ***standard deviation*** (S) for any given variable will be bigger the longer the holding period. The formula for adjusting standard deviation for holding period is:

$$S_{t2} = S_{t1} \times \frac{1}{\sqrt{\frac{t_1}{t_2}}}$$

where t is the holding period.

Example 1: Standard deviation

If the daily standard deviation S1 for a variable is 0.3, what is the 10-day standard deviation S10?

$$S10 = 0.300 \times \frac{1}{\sqrt{\frac{1}{10}}}$$

$$= \frac{0.300}{0.316}$$

= 0.949

(Another way of expressing this is that the standard deviation increases by the square root of the relative length of the time period. So here: 0.300 x √10 = 0.949.)

Standard deviation may also be calculated in absolute terms, in daily change or in percentage terms if the underlying data allows. So for exchange rates we could calculate the daily percentage change in exchange rate and use that to calculate S rather than the absolute value of the exchange rate.

The ***confidence level*** is the probability that a particular level of loss will not be exceeded - typically 95% or 99%. Confidence level is also called confidence interval or confidence limit. In non-financial applications, e.g. quality control of the size of manufactured components, confidence levels usually relate to two-tailed distributions. For VaR, i.e. financial value at risk calculations, the confidence level conventionally relates to a “one tailed" distribution.

Exhibit 1: One and two tailed confidence levels

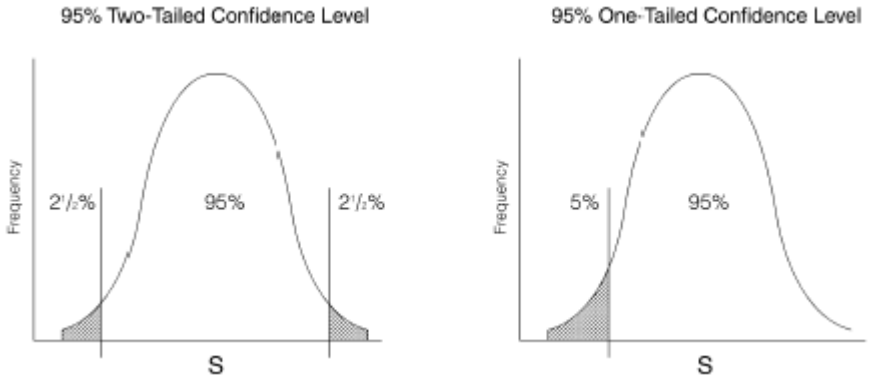

VaR conventionally uses the one-tailed distribution because typically it is concerned only with "loss".

Example 2: Value at Risk

For the 95% one-tailed confidence level how many standard deviations is the 5% point away from the mean?
95% = 0.950
Referring to Appendix 1, 0.950 = 1.65 S

For the company with GBP10m of dollar exposure, what is the daily VaR at 95% confidence level if the standard deviation for daily change in dollar-sterling is 0.300%?

Exposure = GBP 10m against USD
GBP / USD DAILY = 0.300%
95% of the time
Rate will not fall beyond 1.65 x S,
i.e. 1.65 x 0.300% = 0.50%

So VaR = GBP 10m x 0.50%
= GBP 50,000 (rounded)

So, on 95% of days we should expect to lose no more than £50,000 on the exposure – but on 5% we will expect to lose more.

If the holding period for the currency was 10 days, what is the VaR?
We have already converted S1 of 0.300% to S10 of 0.949% (above).
So:
VaR10 = GBP 10m x 1.65 x 0.949
= GBP 156,533, say £157,000

So, over 95% of 10 day periods we should expect to lose no more than £157,000 on the exposure – but on 5% we will expect to lose more...

What is the daily VaR if we work to a 99% confidence level?

99.0% = 0.990
Referring to Appendix 1, 0.490 = 2.33 S
VaR = 10m x 2.33 x 0.300%
= GBP 69,900

So on 99% of days we should expect to lose no more than £69,900 on the exposure – but on 1% of days we will expect to lose more.

2.1 Working with implied volatilities

The above treatment works perfectly for standard deviations derived directly from the observed changes, e.g. percentage changes in daily FX rates. However, volatilities (i.e. standard deviations) are generally quoted in the market on a *lognormal* basis, i.e. they are the standard deviations of the natural logarithm of the percentage change, rather than the standard deviation of the percentage change itself. This is because the Black-Scholes option pricing model is built on lognormal standard deviations.

The procedure for working with such implied volatilities is very similar to working with "direct" volatilities. Taking the example from above:

For the company with GBP10m of dollar exposure, what is the daily VaR at 95% confidence level if the *implied volatility* (i.e. standard deviation) for daily change in dollar-sterling is 0.300%?

Exposure = GBP 10m against USD
GBP / USD S_{DAILY} = 0.300% on an *implied* basis
95% of the time
Rate will not fall beyond 1.65 x S, (as before)
i.e. 1.65 x 0.300% = 0.50% (as before)

This time VaR = GBP 10m x ($e^{0.50\%}$ -1)
= GBP 50,125 (rounded)

Because we are working in lognormal terms, instead of multiplying we use e to the power to calculate the VaR. The difference is small but for larger amounts or higher volatilities it could be significant.

2.2 Applicability

The use of VaR has been spearheaded by banks. Banks have monetary assets as well as monetary liabilities and it is a core part of their business to take and trade positions, which are then managed within money and time limits. So VaR is particularly relevant for banks and VaR methodology has been written into much of the regulatory framework for controlling bank market risk.

For corporates, the technique is less applicable because of point 3 above. In addition, corporates tend to hedge all or most of their financial risks progressively or iteratively based on (notional) principal amounts rather than manage exposures as banks do.

However, the pressure for greater public disclosure of financial risk continues e.g. FRS 13 in the UK and SEC Rules on disclosure in the US. The SEC Rules allow for either tabular or VaR presentation and as VaR reveals less it is becoming more widely used.

One particular extension of this mathematics is to be able to demonstrate graphically the effect of volatility of a risk. The following is a simple application of the principles seen here applied to foreign exchange risk management:

Exhibit 2: Cone of uncertainty

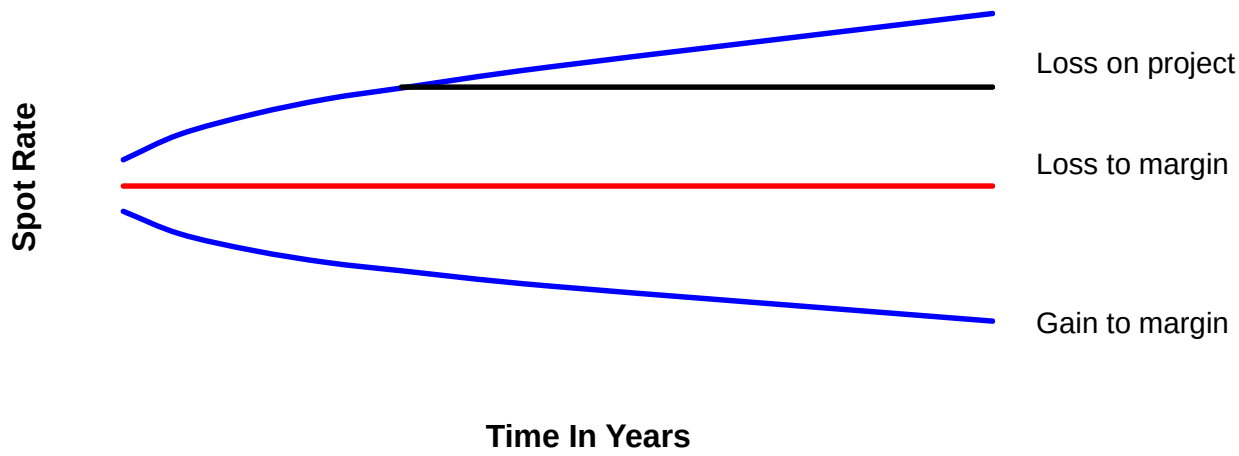

In this case an FX exposure has arisen (selling variable currency proceeds and buying the base) and the current FX rate is 1.5. The curves represent the implied limits of exchange rate movement for a given number of standard deviations from the expected value over the time period shown. Above 1.74, the project moves into loss. The graph illustrates well the concept of implied volatility leading to uncertainty of outcome and the real possibility of its effect on trading situations.

Each point of the graph is calculated as follows:

$$F = S_0 \times e^{Z \times \sigma \times \sqrt{t}}$$

Where:

F is the maximum or minimum expected future spot rate of the variable in question (here the FX rate), at the confidence level implied by Z

S_0 is the current spot rate

Z is the number of standard deviations from the mean, either positive or negative, to provide the required confidence level (e.g. for 95% certainty use Z = 1.65

σ is the (annualised) implied volatility for the period

t is time, in years

This same principle can be applied to portfolios.

3 Portfolio effect

3.1 Concept

Portfolio theory looks at the effect of holding a basket of assets and / or liabilities rather than just the one. At its simplest, the total return or value of a portfolio is equal to the weighted average of the individual returns or values.

Similarly, the degree of fluctuation in value (i.e. risk) of the portfolio will be equal to the weighted average of the individual fluctuations, but only if all portfolio assets move together. For this they have to be perfectly correlated, i.e. all have a correlation coefficient of 1. Now, if the correlation coefficients between elements of the portfolio are less than one (so that fluctuations in the value of the elements don't reinforce each other), then the portfolio risk will be less than the weighted average of individual element risks. (Recall that the correlation coefficient must be between +/-1.)

The benefit of the portfolio approach for treasury is that the cost of putting hedges in place to bring the net portfolio risk within acceptable levels should normally be significantly less than the total cost of hedging each risk individually.

Example 3: Corporate currency risk

European Manufacturer

- Functional currency EUR
- Risk currencies EUR, USD, JPY, AUS
- VAR total, based on 95% confidence limit,
  1 day holing period, for individual currencies
  = EUR 276m
- VAR total, based on 95% confidence limit,
  1 day holding period, on correlated basis
  = EUR 216m
- Diversification Benefit 21.7%

Here we have quantified the risks using VaR. Simply adding the individual VaRs to aggregate the risks is incorrect if they combine into a portfolio and have a correlation coefficient of less than +1. The portfolio basis allows for currency interactions, and indicates that the risk in the portfolio is less than the sum of the components.

By selectively hedging currencies within the portfolio and thus changing their weightings, the VaR of the portfolio can be minimised (but not brought to zero unless all the exposures are zero). Conversely, ill-thought through hedging can actually increase the risk to the firm.

Example 4: Bank market risk
NEDCOR (South Africa) Historical VAR, 99%, 1 day

	Year Average (ZAR '000)
• Foreign exchange	7,564
• Interest rate	13,815
• Equity products	2,744
• Commodities	420
• Individual total	24,543
• Diversification effect	(7,658)
	16,885
• Diversification effect %	31.2%

This shows how VaR can be combined across a portfolio of different assets. Again, the diversification effect arises from the lack of correlation between components. These assets are bank-style financial assets; the values of the FX exposures in the corporate example are the corporate equivalents.

The Value at Risk of a two element portfolio is given by:

$$VaR_{AB} = \sqrt{VaR_A^{\ 2} + VaR_B^{\ 2} + \left(2 \times \rho_{AB} \times VaR_A \times VaR_B\right)}$$

Where:
VaR_A is the VaR for variable A
VaR_B is the VaR for variable B
VaR_{AB} is the correlated VaR for variables A and B
ρ_{AB} is the correlation coefficient between variables A and B

This formula is similar to the calculation of portfolio standard deviation, using the relative weightings of the portfolio elements. This is because VaR is a scaled measure of standard deviation. (The scaling factors being the size of the portfolio, for example £1m, and the number of Standard Deviations for the chosen confidence level, for example 1.65 SDs at 95%.)

We would get the same answer by calculating the portfolio standard deviation directly, and then scaling it back up for the size of the portfolio and the number of SDs.

For a three element portfolio the correlated VaR formula extends to:

VaR_{ABC}
$= \sqrt{[VaR_A^2 + VaR_B^2 + VaR_C^2 + (2 \times \rho_{AB} \times VaR_A \times VaR_B) + (2 \times \rho_{AC} \times VaR_A \times VaR_C) + (2 \times \rho_{BC} \times VaR_B \times VaR_C)]}$

Where:
VaR_A is the VaR for variable A
VaR_B is the VaR for variable B
VaR_C is the VaR for variable C
VaR_{ABC} is the correlated VaR for variables A, B and C
ρ_{AB} is the correlation coefficient between variables A and B
ρ_{AC} is the correlation coefficient between variables A and C
ρ_{BC} is the correlation coefficient between variables B and C

And so on, for larger numbers of portfolio elements. (Start using a computer at this point, if you aren't using one already!)

3.2 Advantages and caveats

You will have noticed that in the examples shown, historical correlated VaR measures of risk are used.

The supreme advantage and attraction of VaR is that because it is measured just in units of money it offers the ability to summarise the whole risk of any organisation into one number at a required confidence level, by treating the entire organisation as one portfolio. The other advantage of VaR is that because it explicitly links loss to probability, it is ideally suited to plotting on a risk map.

For a number of years risk experts have sought to capture this supreme advantage (and in doing so, many have relied on VaR to the exclusion of other measures), but there are some substantial practical issues which reduce its utility.

Practical issues for all users:

- Changes in value may not be normally distributed. For some instruments, e.g. options and indeed foreign exchange, they are definitely not normally distributed.
- Shocks to the system may occur with greater frequency than the normal curve implies – 'fat tails'. So VaR is typically used for normal market conditions and supplemented by "stress testing"/scenario analysis i.e. by modelling the impact of large shocks.
- Frequency distributions, market values and liquid markets exist for commonly used exchange rates and interest rates. They may not for other types of corporate risks.
- For portfolios of risks and particularly for non-financial firms, correlation coefficients may be difficult to establish. You need to generate a correlation coefficient for each pair of variables, and also the future may not be like the past - correlations can be volatile. In the financial crisis of 2008, many assets fell in value together, i.e. their correlations tended towards 1. In such a case the diversification benefit, calculated by VaR, disappears.
- The period used to gather data for historic analysis has to be set with care. A longer period will reflect long-run trends but older data may be outdated; a shorter period will better reflect more recent developments but will be skewed by unusual fluctuations.

  Methods other than the historical variance-covariance method are used to build a more sophisticated approach to volatility and correlation:

Exhibit 3: Alternatives to historic based models

Historic simulation:	Uses historic portfolio returns to avoid the correlation issue.
Monte Carlo simulation:	Uses a large run of scenarios, say 10,000, to generate a distribution, using our own assumptions about future volatility and correlation.
Implied volatility	Volatilities implied by option prices are sometimes used, i.e. the market's best estimate of the volatility to occur over the future period.

- The holding period used in many of the examples above was 24 hours. An appropriate holding period for each risk type needs to be established and corresponding data captured. One key factor here is the degree of liquidity i.e., how fast can the position be liquidated (typically within the holding period). For holding periods exceeding one month, VaR may not be appropriate because statistical reliability tails off.
- Once the portfolio extends beyond two or three assets, computing demands extend dramatically and the model risks becoming hideously complex.

For corporates there are important additional problems:

- Many corporate assets/risk positions cannot be marked-to-market (i.e. accurately valued) nor easily liquidated. Examples are brands, know-how, stock, work-in-progress and perhaps receivables, specialised plant and machinery, or forecast revenue streams.
- There may not be enough data points for statistical reliability.

Note that VaR methodology is now usually available in treasury management systems.