Note
Value-at-Risk and Stress Testing
By A.V. Vedpuriswar

"The difficulty is that VaR takes into consideration only the probability of loss and not the potential size of the loss, lending itself to manipulation."

Phelim Boyle, Mary Hardy, Ton Vorst[1]

Introduction

Risk management has come a long way in the past four decades. Over time, various mathematical models and tools have been developed to quantify and control risk. Complete dependence on models without exercising human judgment and intuition can be disastrous as we have seen during the recent sub prime crisis. At the same time, managing risk based on "philosophy and culture" alone will not take an organization far. In this chapter, we look at probably the most important tool used in risk management, Value-at-Risk and its extensions. The aim is not to create "quants" experts. Rather, it is to familiarize business managers across functions with how Value-at-Risk can be used to measure and control the risks facing a company.

Understanding Value-at-risk

VAR is one of the key building blocks in market risk management. *VAR summarizes the worst loss over a target horizon that will not be exceeded at a given level of confidence*. For example, we may state that, "under normal market conditions, the most the portfolio can lose over a month is about $3.6 million at the 99% confidence level." This means that the 99% monthly VAR is $3.6 million. In simple terms, there is a 99% probability that losses will not exceed $3.6 million during a given month. Or there is only a 1% probability that the losses will exceed $3.6 million.

Jayanth Varma in his book, "Derivatives and Risk Management[2]," has explained in a simple way how VAR can be interpreted in four different ways. Thus 99% VAR can mean the:

a) level of capital that is sufficient to absorb losses 99% of the time
b) level of loss that is exceeded only 1% of the time
c) worst of the best 99% of all possible outcomes
d) best of the worst 1% of possible outcomes

The main idea behind VAR is to get an aggregated rather than a fragmented view of risk. Initially applied to market risk, VAR is now used to measure credit risk, operational risk and even enterprise wide risk. Banks which meet certain norms prescribed under Basle II can use their own VAR models for measuring market risk.

---

[1] The Journal of Derivatives, Fall 2005.
[2] Tata McGraw Hill, 2008.

VAR applications

- **VAR as a benchmark measure**

VAR can be used as a company wide yardstick to compare risks across different markets and businesses over time. VAR can be used to drill down into risk reports to understand whether the higher risk is due to increased volatility in the markets or conscious risk taking.

- **VAR as a potential loss measure**

VAR can give a broad idea of the losses an institution can incur. This in turn can trigger a discussion at the senior levels of management. Are we capable of withstanding such a loss?

- **VAR as an integrated measure of risk**

VAR can be used to integrate all the risks facing the institution - market risk, credit risk, operational risk and other risks.

Exhibit 4.1

One-day, 98% risk management VaR (CHF)

in / end of	Interest rate	Credit spread	Foreign exchange	Commodity	Equity	Diversi-fication benefit	Total
2015 (CHF million)							
Average	20	36	11	2	23	(43)	49
Minimum	6	31	5	1	16	–[1]	34
Maximum	35	42	22	4	35	–[1]	63
End of period	17	40	9	1	31	(42)	56
2014 (CHF million)							
Average	12	32	9	2	18	(31)	42
Minimum	7	28	5	0	13	–[1]	35
Maximum	17	39	17	4	25	–[1]	56
End of period	9	39	7	1	20	(29)	47
2013 (CHF million)							
Average	18	35	9	2	16	(40)	40
Minimum	8	30	3	1	11	–[1]	33
Maximum	45	41	24	4	36	–[1]	55
End of period	10	32	6	3	24	(30)	45

Excludes risks associated with counterparty and own credit exposures.

[1] As the maximum and minimum occur on different days for different risk types, it is not meaningful to calculate a portfolio diversification benefit.

Average one-day, 98% risk management VaR by division (CHF)

in	Swiss Universal Bank	International Wealth Management	Asia Pacific	Global Markets	Strategic Resolution Unit	Diversi-fication benefit[1]	Credit Suisse
2015 (CHF million)							
Total average VaR	6	1	17	53	17	(45)	49
2014 (CHF million)							
Total average VaR	2	0	11	45	13	(29)	42
2013 (CHF million)							
Total average VaR	2	0	11	41	16	(30)	40

Excludes risks associated with counterparty and own credit exposures. Investment Banking & Capital Markets has only banking book positions. The calculation of divisional average risk management VaR under the new organization required certain additional assumptions and allocation methods, which may not be required for future periods given the level of information then available.

[1] Difference between the sum of the standalone VaR for each division and the VaR for the Group.

Source: Credit Suisse Annual Report, 2015

Exhibit 4.2

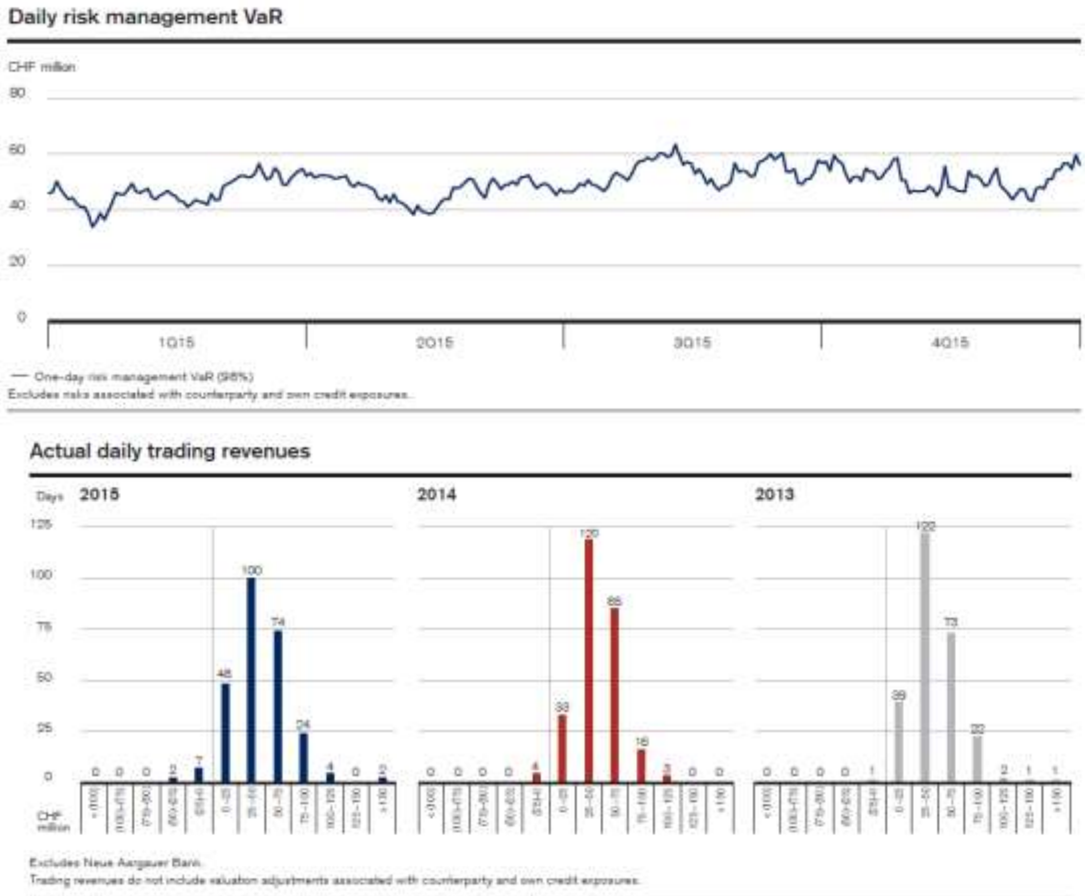

Source: Credit Suisse Annual Report, 2015

- **VAR as an information reporting tool**

VAR is a useful information reporting tool that facilitates disclosure of aggregated risk without revealing individual positions. Nearly all large financial institutions report quantitative information about market risk using VAR. Many of them provide summary VAR figures on a daily, weekly or monthly basis. Such disclosures are an effective means of enforcing market discipline.

- **VAR as a risk control tool**

VAR is a useful risk control tool that can supplement position limits. In volatile environments, VAR can be used as the basis for scaling down positions. Here VAR scores over position limits. In addition, VAR accounts for diversification, unlike position limits.

- **VAR as a measure of economic capital**

VAR can be viewed as a measure of risk capital or economic capital, the aggregate capital required as a cushion against unexpected losses. Banks routinely calculate economic capital using a high confidence level, eg: 99.98%.

- **VAR as a risk adjusted performance measure**

  VAR can be used to arrive at a risk adjusted performance measure. Without controlling for risk, traders may become reckless. The compensation received by traders has an asymmetric pay off profile. When traders make a large profit, they receive a huge bonus. When they make a loss, the worst that can happen is they will get fired. The pay off profile is similar to that of a long position in a call option, i.e., unlimited upside but limited downside. Risk adjusted compensation can help in curbing the temptation to indulge in reckless behaviour.

- **VAR as a Strategic tool**

VAR can be used as a strategic tool by top management to identify where shareholder value is being added. This can facilitate better decisions about which business lines to expand, maintain or reduce. Executives are forced to examine prospects for revenues, costs and risks in all their business activities. As managers start to learn new things about their business, the general quality of management improves and there is better capital deployment.

Exhibit 4.4
VAR applications

Passive role	Reporting risk	• Disclosure to share holders • Management reports • Regulatory requirements
Defensive role	Controlling risk	• Setting risk limits
Active role	Allocating risk	• Performance evaluation • Capital allocation • Strategic business decisions

Based on the work of Philippe Jorion.

- **VAR & investment management**

VAR is becoming more relevant to the investment management industry, both in asset allocation and Fund Management. Using VAR systems, investors can monitor their market risk better. Passive asset allocation or benchmarking, does not keep risk constant because the composition of the indices can change substantially. VAR can identify such trends. VAR tools are also useful in allocating funds across asset classes.

Active portfolio management may also change the risk profile of the fund. A sudden increase in the reported VAR should prompt a deeper analysis of the situation. Is more risk being taken? Are unauthorized trades being made? Is the risk increase justified by current conditions? Are different managers making similar bets? Different investment managers, acting in isolation, may be simultaneously increasing their exposure to a sector which is looking attractive. So a centralised VAR system can identify such trends and facilitate corrective action, if required.

- **VAR & risk budgeting**

VAR can also facilitate risk budgeting, a concept that is becoming popular in investment management. Risk budgeting essentially means a top down allocation of economic risk capital starting from the asset classes down to the choice of the active manager and even to the level of individual securities.

VAR Computation

In simple terms, in computing VAR, we first understand the various risk factors that may influence the value of the portfolio. Then we compute the value of the portfolio under various scenarios. Alternatively, we can examine how the portfolio has behaved historically. We study the distribution of the portfolio returns and determine what is the maximum loss likely to be, at a given confidence level. We can do this either by using a simple percentile approach or by using a statistical distribution. We shall examine these methods in more detail a little later in the chapter.

The starting point in VAR is identifying the various risk factors and how they affect the different instruments. If the portfolio consists of a large number of instruments, it would be too complex to model each instrument separately. The first step is mapping. Instruments are replaced by positions on a limited number of risk factors. This simplifies the calculation significantly.

Two broad approaches to valuing the instruments are available. The more straight forward *Local valuation methods* make use of the valuation of the instruments at the current point and incrementally as we move away from the point using the first and perhaps, the second partial derivatives. The entire portfolio is valued only once. The value at other points is calculated by adjusting the base or anchor value suitably. Such an adjustment can normally be made in two ways:

The *delta normal method* assumes that the portfolio measures are linear and the risk factors are jointly normally distributed. Delta is nothing but the rate of change in portfolio value with respect to the underlying asset price. In such cases, daily VAR is adjusted to other periods, by scaling by a square root of time factor. This adjustment assumes that the daily returns are independently and identically distributed. So the variances can be added. The delta normal method is computationally fast even with a large number of assets because it replaces each position by its linear exposure. This method is not appropriate when there are fat tails in the distribution and non linear instruments exist in the portfolio. The delta normal approach can be represented by the equation: $dp = \Delta ds$. Where dp is change in portfolio value, ds is change in underlying price.

If we have the following data, it is a simple matter to calculate the VAR:

- Size of the position
- Volatility of daily returns

- Confidence level
- Time horizon

If we take the average return of the portfolio as the reference point, then VAR is nothing but the product of the position size, the Z value (the distance from the mean in terms of standard deviations,) volatility (standard deviation of daily returns), and the square root of time. We shall discuss how to estimate volatility in the next chapter.

Illustration

Consider an asset valued at $1 million with volatility of daily returns being 10%. What is the daily VAR at 95% confidence level? What will be the 10 day VAR?

From normal distribution tables, we read out the value of Z as 1.645. Note that we are applying a left tail situation as we are concerned about the downside, not the upside.

VAR = (1) (.10) (1.645) = $.1645 million
= $ 164,500

To calculate the 10 day VAR we have to scale by square root of time.

So 10 day VAR = (164,500) √10 = $520,195

Illustration

The 10 day 99% regulatory VAR for UBS as given in Exhibit 5.1 is SF 485 million as on 31 Dec 2008. What would be the 95% 10 day VAR? What would be the 99% daily VAR?

Z value for 95% confidence level is 1.645 while that for 99% confidence level is 2.33

So 95% VAR = $\frac{1.645}{2.33} \times 485$ = SF 342.41 million

1 day 99% VAR = $\frac{485}{\sqrt{10}}$ = SF 153.37 million

When two assets are combined, the volatility of the portfolio has to be computed, using the well known formula, $\sigma = \sqrt{\sigma_1^{\,2} + \sigma_2^{\,2} + 2\rho\sigma_1\sigma_2}$ where σ_1, σ_2 represent the volatility of individual assets and σ that of the portfolio.

This formula can be adjusted as the number of assets increases. If there are three assets, the portfolio standard deviation

$$\sigma = \sqrt{\sigma_1^{\,2} + \sigma_2^{\,2} + \sigma_3^{\,2} + 2\rho_{12}\sigma_1\sigma_2 + 2\rho_{13}\sigma_1\sigma_3 + 2\rho_{23}\sigma_2\sigma_3}$$

The *delta gamma* method incorporates a second order correction to the delta normal VAR by using gamma. Gamma,γ is nothing but the rate of change in delta with respect to the underlying spot price. Long positions in options with a positive gamma have less risk than that implied by a linear model, while short positions in options have greater risk.

The delta gamma approach can be represented by the following equation:

$$dp = \Delta ds + ½ \gamma (ds)^2$$

Here dp is the change in portfolio value, $\Delta = \frac{dp}{ds}$ $\gamma = \frac{d^2 p}{ds^2}$

For more complex pay offs, local valuation is not enough. The entire portfolio must be revalued at different points instead of making adjustments to an anchor value. Take the case of a long straddle, i.e, the purchase of call and a put with the same strike price. The worst pay off (sum of the two premiums) will be realized if the spot rate does not move at all. So if we value the portfolio at a few extreme points, we will not get the full picture. All intermediate values must be checked. This is where *full valuation methods* come in handy. These methods reprice the instruments over a broad range of values for the risk factors. Two popular full valuation methods are Historical Simulation and Monte Carlo simulation. We will now examine them in detail.

Historical simulation

The historical simulation method consists of going back in time and examining past data. Many global banks use five years of past data. If a year contains 260 trading days, it means 1300 daily returns on the portfolio are tabulated in ascending order. This method makes no specific assumption about the return distribution. All it needs is historical data. This is an improvement over the normal distribution because historical data typically contain fat tails. Essentially, historical simulation applies the percentile method of calculating dispersion. The past returns on the portfolio are tabulated in ascending order. Depending on the confidence level, the bottom 5% or 1% or .1%, are marked off to get the VAR estimate. The traditional school of thought advocated going far back in time to obtain adequate data from which meaningful inferences can be made. But the problem here is that this may involve observations that are no longer relevant. Indeed longer sampling paths can mask current trends when a sudden change in the market environment occurs. This was so during the sub prime crisis. A small example will illustrate how historical simulation is done.

Illustration[3]

% Returns	Frequency	Cumulative Frequency
- 14	1	1
- 12	1	2
- 10	1	3
- 8	2	5
- 5	1	6
- 4	3	9
- 3	1	10
- 1	2	12
0	3	15
1	1	16
2	2	18
3	1	19
4	1	20
5	1	21
6	1	22
7	1	23
8	1	24
9	1	26
11	2	27
13	1	28
15	1	29
20	1	30

What is VAR (90%)?

There are 30 observations. These are already arranged in ascending order. 10% of 30 is 3. We notice by inspection that 3 observations lie below – 8. So VAR is – 8. Of course 30 is too small a number. To get a meaningfully accurate VAR estimate, we would need a much larger number of observations.

Monte Carlo Simulation

The most advanced and sophisticated VAR modeling technique is Monte Carlo Simulation. This method does not use historical data. A probability distribution is specified for the random variable based on a good understanding of its past behaviour.

[3] Adapted from Philippe Jorion, "Financial Risk Manager Handbook," John Wiley & Sons, 2007.

Then using random numbers, the portfolio returns are simulated. From these returns, VAR is estimated.

The Monte Carlo method can incorporate a wide range of risks including price risk, volatility risk, fat tails and extreme scenarios. Non linear exposures and complex pricing patterns can also be handled. Monte Carlo analysis can deal with time decay of options, daily settlements and associated cash flows and the effect of pre specified trading or hedging strategies.

Different random numbers will lead to different results. So a large number of iterations may be needed to converge to a stable VAR measure. The Monte Carlo approach is thus computationally quite demanding. The method may also take far too much time even with the best of computing resources.

To speed up the computation, various methods have been devised. In the *Grid Monte Carlo approach*, the portfolio is exactly valued over a limited number of grid points. For each simulation, the portfolio is valued using a linear interpolation from the exact values at adjoining grid points. Sometimes, the simulation can be speeded up by sampling along the paths that are most important to the problem at hand. For example, if the goal is to measure a tail quantile accurately, there is no point in doing simulations that will generate observations in the centre of the distribution.

To increase the accuracy of the VAR estimator, we can partition the simulation region into two or more zones. An appropriate number of observations is drawn from each region. Using more information about the portfolio distribution results in more efficient simulations. The simulation can be done in two phases. The first pass runs a traditional Monte Carlo. The risk manager then examines the region of the risk factors that cause losses around VAR. A second pass is then performed with many more samples from the region.

The accuracy of the results and the predictive power of a Monte Carlo simulation will be as good as the model underlying it. The Monte Carlo approach requires users to make assumptions about the stochastic process and to understand the sensitivity of the results to these assumptions. Indeed, the first and most crucial step of Monte Carlo simulation consists of choosing a particular stochastic model for the behavior of prices.

Selecting the right probability distribution calls for a good understanding of the market variable. For example, the geometric Brownian motion model adequately describes the behaviour of stock prices and exchange rates but not that of fixed income securities. In Brownian motion models, price shocks are never reversed and prices move as a random walk. This is not an appropriate price process for default free bond prices which are characterized by mean reversion and must converge to their face value at expiration.

Choosing the time horizon

What is the most appropriate measure: 1 day VAR, 10 day VAR, monthly VAR or annual VAR? The longer the time horizon, the greater the VAR measure. As we have seen earlier, volatility is proportional to the square root of time. So the scaling is proportional to the square root of the time horizon. The horizon should be chosen depending on the context. If a portfolio can be modified or rebalanced quickly, a daily VAR is appropriate. In such a situation, increasing the time horizon does not make sense. On the other hand, if the portfolio is "sticky" and cannot be rebalanced quickly, then a longer time horizon makes sense. Loan portfolios for example are not marked-to-market. Moreover, problems in case of loans may not appear immediately. So where credit risk is the dominating factor, a longer time horizon is preferable. Similarly for illiquid portfolios, a longer time horizon is recommended.

VAR, as mentioned earlier, can also be used as a measure of economic capital. If the VAR number is being used to decide how much capital is to be set aside to avoid bankruptcy, a long time horizon is advisable. Institutions will need time for corrective steps when problems start to develop. Indeed, raising capital itself, can be a challenge during a crisis. The longer the time needed for corrective action, the more the capital that must be set aside. By increasing the time horizon, the bank can be more conservative with respect to its capital needs.

Backtesting

Models must be tested from time to time to check whether they are functioning well. Backtesting is the process of using actual market data to check the accuracy of the model. It helps in verifying whether actual losses are in line with projected losses. A rigorous back testing process forms the underpinning for the Basle framework. If the model is robust, back testing will reveal only a few VAR exceptions, probably a couple of exceptions in a year.

The Basle framework has stipulated that if a 99% VAR is backtested with 250 days of data, the results can be divided into green, yellow and red zones. In the green zone, the 99% VAR is exceeded less than 5 times. In the yellow zone, there may be 5 or more exceptions but less than 10. The model could be underestimating VAR in this case but the evidence is not conclusive. In the red zone, there are 10 or more exceptions. This is a clear indication that the model is not working satisfactorily.

The existence of clusters of exceptions indicates that something is wrong. Credit Suisse reported 11 exceptions at the 99% confidence level in the third quarter of 2007, the erstwhile Lehman brothers three at 95%, Goldman Sachs five at 95%, Morgan Stanley six at 95%, the erstwhile Bear Stearns 10 at 99% and UBS 16 at 99%. Clearly, VAR is a tool for normal markets and it is not designed for stress situations.

What should be the time horizon and confidence level for a backtest? Backtesting must strike a balance between two types of errors: rejecting a correct model vs accepting an incorrect model. For example, too high a confidence level may reduce the expected number of observations in the tail and thus the power of the tests. So backtesting is often done with a somewhat lower confidence level.

Illustration

Based on a 90% confidence level, how many exceptions in back testing a VAR should be expected over a 250 day trading year?

Since the confidence interval is 90%, 10% of the time, loss may exceed VAR

So no. of exceptions = (.10) (250) = 25

Illustration

Let us say we back test a model using 600 days of data.
The VAR confidence level is 99% and there are 9 exceptions.
Should we reject the model?
For each observation,

Probability of exception = .01

Probability of no exception = .99

Let us first find the probability of having 8 or fewer exceptions.

Probability of no exception in the data set	=	$(.99)^{600}$	= .0024
Probability of 1 exception in the data set	=	$600C_1\,(.99)^{599}\,(.01)$	= .0146
Probability of 2 exceptions in the data set	=	$600C_2\,(.99)^{598}\,(.01)^2$	= .0441
Probability of 3 exceptions in the data set	=	$600C_3\,(.99)^{597}\,(.01)^3$	= .0888
Probability of 4 exceptions in the data set	=	$600C_4\,(.99)^{596}\,(.01)^4$	= .1338
Probability of 5 exceptions in the data set	=	$600C_5\,(.99)^{595}\,(.01)^5$	= .1612
Probability of 6 exceptions in the data set	=	$600C_6\,(.99)^{594}\,(.01)^6$	= .1614
Probability of 7 exceptions in the data set	=	$600C_7\,(.99)^{593}\,(.01)^7$	= .1384
Probability of 8 exceptions in the data set	=	$600C_8\,(.99)^{592}\,(.01)^8$	= .1036

Cumulative probability of getting 8 or fewer exceptions

= .0024 + .0146 + .0441 + .0888 + .1338 + .1612 + .1614 + .1384 + .1036

= .8483

Probability of getting 9 or more exceptions = 1 - .8483 = .1517

≈ 15.2%

> 5%

So if we are testing the hypothesis at 5% confidence level, we cannot reject the model.

Suppose we get 10 exceptions

Expected probability = $1 - [.8483 + 600C_9\,(.99)^{591}\,(.01)^9]$

= 1 - .9171 = .0829 ≈ 8.3%

> 5%

Again we cannot reject the model.

Suppose we get 11 exceptions

Expected probability = $1 - [.9171 + 600C_{10}\,(.99)^{590}\,(.01)^{10}]$

= 1 – [.9171 + .0411]

= 1 - .9582 = 4.18%

< 5%

So the model should be rejected if we get 11 or more exceptions.

Illustration

A 99% VAR model reports 5 exceptions in a year.

Assuming there were 250 trading days, test at a 5% significance level whether the model must be rejected. If the model reported 6 exceptions, would your recommendation change?

Solution

Probability of no exception	=	$(.99)^{250}$	=	.0811
Probability of 1 exception	=	$250C_1(.99)^{249}(.01)$	=	.2047
Probability of 2 exception	=	$250C_2(.99)^{248}(.01)^2$	=	.2574
Probability of 3 exception	=	$250C_3(.99)^{247}(.01)^3$	=	.2149
Probability of 4 exception	=	$250C_4(.99)^{246}(.01)^4$	=	.1341
Probability of having 4 or less exceptions	=			.8922
Probability of having 5 or more exceptions	=	1 - .8922	=	.1078
			=	10.78%
			>	5%

So at a 5% significance level, the model cannot be rejected.

Probability of having 6 or more exceptions	=	$1-[.8922+250C_5(.99)^{245}(.01)^5]$
	=	1 – [.8922 + .0666]
	=	.0412
	=	4.12%
	<	5%

So at a 5% significance level, the model must be rejected if there are 6 exceptions.

Illustration

Suppose we have to scale volatility from one day to 5 days given an autocorrelation of 0.2 between the daily returns of successive days. How will you do this? Assume the daily volatility, σ does not change.

Let the daily volatility be 1.
Then the daily variance is also 1.

Autocorrelation means each day's price movements will have an impact on the next day's movements through the correlation coefficient. On the other hand, today's price movement will have an impact on the price movement two days later through the square of the correlation coefficient and so on.

So variance for 5 days

$= \sigma^2 + \sigma^2 + \sigma^2 + \sigma^2 + \sigma^2$
$+ (2)(\rho)(\sigma^2) + (2)(\rho^2)(\sigma^2) + (2)(\rho^3)(\sigma^2) + (2)(\rho^4)(\sigma^2)$
$+ (2)(\rho)(\sigma^2) + (2)(\rho^2)(\sigma^2) + (2)(\rho^3)(\sigma^2)$
$+ (2)(\rho)(\sigma^2) + (2)(\rho^2)(\sigma^2)$
$+ (2)(\rho)(\sigma^2)$

$= 5\sigma^2 + 8\rho\sigma^2 + 6\rho^2\sigma^2 + 4\rho^3\sigma^2 + 2\rho^4\sigma^2$

So the variance for 5 days

$= (5)(1)^2 + (8)(.2)(1)^2 + (6)(.2)^2(1)^2 + (4)(.2)^3(1)^2 + (2)(.2)^4(1)^2$
$= 5 + 1.6 + .24 + .032 + .0032$
$= 6.8752$

So 3 day volatility $= \sqrt{6.8752} = 2.62$

Problem

The volatility of daily returns does not change but there is auto correlation among the daily returns on different days. We want to scale from one day VAR to 5 day VAR. How can we do this if the correlation between two succeeding days is 0.1?

Solution

Assume the daily volatility is 1 and the daily variance is also 1. The required variance for 5 days

$= 5\sigma^2 + 8\rho\sigma^2 + 6\rho^2\sigma^2 + 4\rho^3\sigma^2 + 2\rho^4\sigma^2$
$= 5 + (8)(.1) + (6)(.1)^2 + (4)(.1)^3 + (2)(.1)^4$
$= 5 + .8 + .06 + .004 + .0002 = 5.8642$

5 day volatility $= \sqrt{5.8642} \approx 2.42$

Effectiveness of VAR models

How effective are VAR models? Do they hold good in real life? In the wake of the sub prime crisis, many economists and experts have held VAR models responsible for the failure of risk management systems. Indeed VAR models completely failed in estimating the potential losses arising out of exposure to sub prime mortgage securities, as the environment suddenly transitioned from a period of relatively benign volatility into a highly volatile one.

A 99.7% confidence interval, corresponding to a 3σ test should under ordinary circumstances give us a fairly good handle on the risks involved. But then the sub prime crisis could hardly be described as an "ordinary" set of events. Indeed Goldman Sachs' chief financial officer David Viniar once described the credit crunch as "a 25-sigma event". One could obviously not expect a '3σ' model to be appropriate for a '25σ' situation! The fact that such events can happen is a clear indication that real life probability distributions have fat tails. The normal distribution is not meanigful for such black swans.

Some VAR models failed during the sub prime crisis because they used historical simulation based on five years of historical data. A longer observation period can smoothen out business cycles and incorporate a wider variety of market conditions. But the problem with such models is they do not react quickly to a sudden change in circumstances. The historical data pertaining to the period 2002-2006 completely masked the paradigm shift that took place starting in the second quarter of 2007. Indeed, the type of VAR model that would actually have worked best in the second half of 2007 would have had a frequently updated short data history, that weighted recent observations far more heavily than distant observations.

To make VAR models more responsive, monthly or even quarterly updating of the data series may become the norm. Shifting to weekly or even daily updating would improve the responsiveness of the model to a sudden change in market conditions.

Despite its limitations, the utility of VAR cannot be questioned. VAR methods represent the culmination of a trend towards centralized and integrated risk management. Such a "global" approach to measuring risk makes sense because the sources of risk have multiplied and volatility has increased. A portfolio approach, gives a better picture of risk, compared to a piecemeal view of different instruments. The essence of VAR is a portfolio approach. Ofcourse VAR must be complemented by stress testing. We shall examine stress testing shortly.

Conditional VAR

The problem with VAR is that it tells us about the losses within a reasonable confidence interval. But what if we are outside the confidence interval, i.e., we are in the tail of the distribution. The conditional VAR is a more useful measure in this context. Conditional VAR takes into account the losses when things get out of control. As a concept, conditional VAR is more difficult to understand. It is also difficult to back test. However, conditional VAR can be a useful supplement to VAR. It can be a better basis for working out risk adjusted compensation for traders. The conditional VAR pays greater importance to the "black swan." The conditional VAR is also called the expected shortfall or tail loss. However, conditional VAR must not be confused with Extreme Value Theory or Stress Testing.

Illustration[4]

Consider an investment with 99.1% VAR being \$1 million. The probability of a loss of \$10 million has been estimated at .9%. What is the conditional VAR at a 99% confidence level?

We can say that in the 1% tail region, the probability of loss of \$1 million is $\frac{.1}{1} = 10\%$

and the probability of loss of \$10 million is $\frac{.9}{1} = 90\%$

So the conditional VAR/expected shortfall = (.1) (1) + (.9) (10)
= \$ 9.1 million

Extreme Value Theory (EVT)

The normal distribution does not give too much importance to the tails. When we are sure of 99.7% of the value being within a bound, why bother about the remaining .3%? Unfortunately, as past financial crises have amply indicated, tail events can have catastrophic consequences. So we must treat the tails of the distribution with abundant caution. When dealing with exceptional events which stretch far into the tails than a 3σ situation, we need other techniques to supplement VAR. Extreme Value Theory (EVT) comes in handy here.

EVT extends the central limit theorem which deals with the distribution of the average of identically and independently distributed variables from an unknown distribution to the distribution of their tails. The EVT approach is useful for estimating the tail probabilities of extreme events. EVT helps us to draw smooth curves through the extreme tails of the distribution.

EVT has been used in the risk assessment of catastrophic events. The impetus for EVT came from the collapse of sea dikes in the Netherlands in February 1953, which led to the flooding of large parts of the country. EVT led to the design of a dike system that could withstand a 1250 year storm.

[4] John. C. Hull, "Risk Management and Financial Institutions," Pearson Education, 2007.

Stress Testing[5]

Whether a 10+ sigma event happens or not, it makes sense to be prepared for the same. Stress testing involves examining the performance of the portfolio under extreme market moves, that can cause extraordinary losses. Early stress testing techniques involved sequentially moving key variables by a large amount. These techniques were essentially some form of sensitivity analysis. Now they have become more sophisticated. Some scenarios are historical, others prospective. Stress tests, which represent abnormal scenarios can be combined with VAR (which corresponds to normal scenarios), to arrive at risk capital. When stress tests indicate large potential losses, either the capital can be enhanced or the exposure can be reduced.

Construction of scenarios is as much art as science. But a few guiding principles can be outlined. Scenarios can be event driven or portfolio driven. In the first case, scenarios can be formulated from plausible events that generate movements in the risk factors. In the second case, the risk vulnerabilities in the portfolio that lead to adverse movements in risk factors are first identified. Some of the well known historical scenarios include:

- The 1987 US stock market crash.
- The 1992 European Monetary System Crisis
- The 1997 Asian currency crisis
- The 1998 LTCM/Russian crisis

During a day, a fall in the equity index by more than 10%, a fall in a globally traded currency by more than 5% against another globally traded currency or an interest rate cut of more than 200 basis points by a central bank are good examples of stress scenarios.

To conclude this section, stress testing is a good complement to VAR because it is difficult to take into account extreme events with the probability distributions with which we are familiar. A five standard deviation move in a market variable would seem unlikely if we go by a standard normal distribution. But in real life, five sigma moves are not uncommon. For example, on October 19, 1987, the S&P 500 moved by 22 standard deviations. On April 10, 1992, 10 year bond yields moved by 8.7 standard deviations. While visualizing such situations and their potential impact, stress testing comes in handy.

[5] Read the article by P Kupiec, "Stress testing in a Value at risk framework," Journal of Derivatives, 6 (1999) pp. 7-24.

Managing model risk[6]

Many VAR models failed during the sub prime crisis. So it may not be out of place here to cover briefly the subject of model risk. Models specify a relationship between various inputs and based on this relationship compute the required output. Some models are fundamental and start from first principles to establish a relationship between various input and output variables. In contrast, statistical models use data to determine correlations without any attempt to find causal relationships. Whatever be its technical characteristics, a model is a simplified representation of reality. Various assumptions are made while constructing a model. If these assumptions are violated, the model will naturally not hold good.

Model risk is not a huge issue when we are dealing with simple, linear instruments. But for exotic derivative instruments, model risk can be high because of interactions between risk factors, lack of transparency, etc. In risk management, often attempts are made to aggregate various positions. In such situations, models may give erroneous results.

In general model risk may arise out of:

- Incorrect model specification
- Incorrect model application
- Faulty implementation
- Incorrect calibration
- Programming issues
- Poor quality of data
- Behavioral issues

Models may not be specified correctly due to various reasons. The model may be using a wrong stochastic process, say a normal distribution, when the tails are "fat." Some risk factors might have been ignored while developing the model. Relationships might have been misspecified. Some models also make very simplistic assumptions such as ready liquidity and zero transaction costs.

Risks also arise when models are incorrectly applied. A wrong model might be in use. Alternatively, the model might not have been updated. For example, models based on 5 year historical simulation would not have worked effectively at the start of the sub prime crisis when market conditions began to change dramatically. In some cases like Monte Carlo simulation, the problem might be too few runs or a poor random number generator.

---

[6] This part draws heavily from the book, "Measuring market risk," by Kevin Dowd, John Wiley & Sons, 2005.

Models may also be implemented in a faulty manner. No model can specify the method of implementation for all possible scenarios. Decisions need to be made by users with respect to valuation, mapping, etc. If these decisions are wrong, the implementation will be faulty. Incorrect calibration is another source of market risk. Calibration problems are particularly common in case of volatility and correlation. The true volatility can be more or less than the estimated volatility. During a crisis, correlations tend to move towards 1. The model might over estimate the benefits of diversification. Problems could also arise on account of the software. There might be bugs in the program. When programs are revised by people who did not originally write them, there could again be problems.

Data problems also contribute to model risk. How we handle time (Calendar time, trading time), and construct data (actual traded data or end of day data) can make a big difference in the way the model works.

Behavioral issues can also contribute to model risk. Traders often have a good understanding of the errors made while estimating the parameters used in the model. They often know which positions understate risks and which overstate them. If traders receive risk adjusted compensation, they will seek out positions with downward biased VAR estimates. There is a strong incentive for traders to 'game' the system.

To deal with model risk calls for a fundamental understanding of the markets and the instruments involved. Senior managers must have a good understanding of the issues involved, so that they can understand the language of the risk managers.

While developing the model, care should be taken to separate the exogenous (causal) variables from the endogenous (caused) variables. Care should also be exercised to distinguish between measurable and non measurable variables. Then a call should be taken on whether a proxy can be found for the non measurable variable. Alternatively, the non measurable variable must be implicitly solved from other variables.

All the assumptions made while developing the model must be carefully evaluated on an ongoing basis. The models must be tested on simple problems with known solutions to see if there is any unexpected response. Backtesting and stress testing must also be done periodically to check for exceptions. Even small discrepancies should not be brushed aside. Such discrepancies may be an early warning of things likely to go wrong in a big way at a later date.

Senior managers must be aware that when a model which performs well in some situations is extended to other situations, the model may fail. Another point to keep in mind is that the utility of a model may not be enduring. If more traders start swearing by the same model and as a result, pursue similar trading strategies, the initial profits will rapidly fall. Some firms which rush into the market before understanding the pros and cons may end up making losses.

Senior managers can also encourage a multidisciplinary approach to team building. Such a team should include mathematicians, computer scientists, finance experts, accounting professionals and model users. Diversity of views and constructive criticism by different stakeholders can ensure that mistakes at the model specification stage are reduced, if not eliminated.

All risk models must be carefully documented. The maths involved, the components, computer code and methods of implementation should be recorded carefully. This will enable risk managers to examine and validate the models on an ongoing basis. The middle office should have enough information to be able to check the model or model results at any time. Risk managers must be able to access the log of model performance with a special focus on any problems encountered and how they have been addressed. The independent middle office should have a clear mandate and authority to block any inappropriate trading or asset management activity in the bank and have authority over the use of pricing/risk models. The middle office should take charge of stress testing, back testing and contingency planning to ensure that all models are adequate for the tasks being handled.

Conclusion

Value-at-Risk is one of the staple tools of modern day financial risk management. It is a simple and elegant way of understanding the maximum losses to a portfolio at a given confidence level and time horizon and accordingly arriving at the amount of capital backing that is needed. But in this chapter, we have also seen the limitations of VAR. VAR models must not be blindly applied. Moreover, VAR must be complemented by other tools such as stress testing to arrive at a true understanding of the risk situation facing a financial institution.

References:

- D. Duffie, J. Pan, “An Overview of value at Risk,” Working Paper, Stanford University, 1997.
- P Kupiec, “Stress testing in a Value at risk framework,” Journal of Derivatives, 6 (1999) pp. 7-24.
- Kevin Dowd, “Measuring market risk,” John Wiley & Sons, 2005.
- Phelim Boyle, Mary Hardy, Ton Vorst, “Life after VaR,” The Journal of Derivatives, Fall 2005.
- John C Hull, “Options, Futures and Other Derivatives,” Prentice Hall, 2006.
- Philippe Jorion, “Financial Risk Manager Handbook,” John Wiley & Sons, 2007.
- Aswath Damodaran, “Strategic Risk Taking – A Framework for Risk Management,” Wharton School Publishing, 2007.
- Philippe Jorion, “Value at Risk: The New Benchmark for Managing Financial Risk,” McGraw Hill, 2007.

- John C Hull, “Risk Management and Financial Institutions,” Pearson Education, 2007.
- Jayanth Rama Varma, “Derivatives and Risk Management” Tata McGraw Hill, 2008.
- UBS Annual Reports.
- Deutsche Bank Annual Reports.
- Goldman Sachs Annual Reports.
- Credit Suisse Annual Reports.
- HSBC Annual Reports.