ISMA CENTRE - THE BUSINESS SCHOOL
OF THE FINANCIAL MARKETS

UNIVERSITY OF READING
ENGLAND

IFID Certificate Programme

Portfolio and Risk Management

Risk Management

Table of contents

1. Overview

In module Portfolio Management we saw how the risks on financial instruments net out and interact with each other when combined in investment or trading portfolios.

Portfolio risk is also the concern of risk managers in investment banks and we conclude this module – indeed the IFID Certificate Programme - with a short introduction to the methodology and the issues surrounding the calculation of **value at risk** on dealing room positions.

Learning Objectives

By the end of this module, you will be able to:

1. Calculate the simple value at risk (VAR) on a bond portfolio, given the return volatilities on the constituent assets and a z score
2. Calculate the diversified value at risk (VAR) on a bond portfolio, given the return volatilities on the individual assets, their correlations and a z score
3. Apply the square root of time rule to calculate the VAR on a position over a given period, given its daily VAR
4. Identify some of the limitations of parametric VAR calculations as a methodology for estimating risk, in particular credit risk
5. Describe in outline the following alternative techniques for estimating the VAR on a portfolio:
    - Historic simulation
    - Monte Carlo simulation
    - Stress testing
6. Outline the procedure for calculating market risk using internal models permitted under the *1996 Amendment* to the Basle Accord
7. Explain the importance of back-testing in validating the internal VAR models

2. Measuring Risk

Derivatives need to be well controlled and understood but we believe that we do that here. Baring Brothers 1993 Annual Report.

As we have seen in this programme, standard deviation of percentage returns is the unit of risk that is common to portfolio management as well as to the options market.

The figure below illustrates the two steps involved in going from historic changes in cash prices or cash returns (in this case on a 10-year French zero) into a **histogram** of percentage returns, from which we can calculate the standard deviation of percentage returns[1].

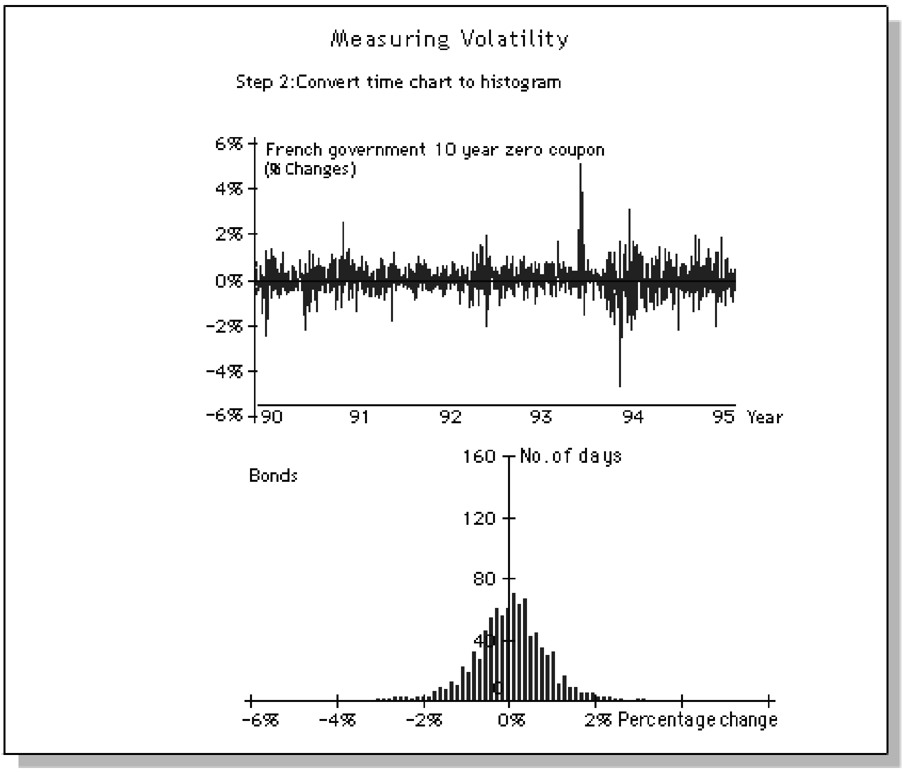

Standard deviation is also the unit of market risk[2] used in a bank's dealing room and in this final section of the programme we take a brief look at how risk that is expressed as a standard deviation can be converted back into cash equivalent amounts for the purpose of ensuring that a bank has enough capital to cover the risks on its dealing operations.

[1] For convenience throughout this section we shall equate the percentage price changes on assets with their total return, but strictly speaking of course the two are not identical.

For a refresher on the concept of a histogram and how to calculate a standard deviation, please consult the background materials on statistics provided in the Pre-requisites - Background Materials section of this training programme.

[2] For a refresher on the normal distribution, please consult the background materials on statistics provided in the Pre-requisites - Background Materials section of this training programme.

3. Value at Risk

3.1. Simple VAR

Value at risk (**VAR**): a measure of the cash loss that could be incurred on a position in a financial instrument, with a given probability, as a result of a adverse changes in the market price of that instrument.

We can easily convert a standard deviation of an asset's historic perentage price into a VAR measure if we make one important assumption: that the observed variability of the return on the asset around its historic average is normally distributed[3].

Suppose the French bond market example on the previous page has a daily average return (μ) of +0.1% and a standard deviation (σ) of 0.55% during the period under review. In the figure below we fit a normal distribution on the histogram of this market's percentage daily price changes that we developed earlier.

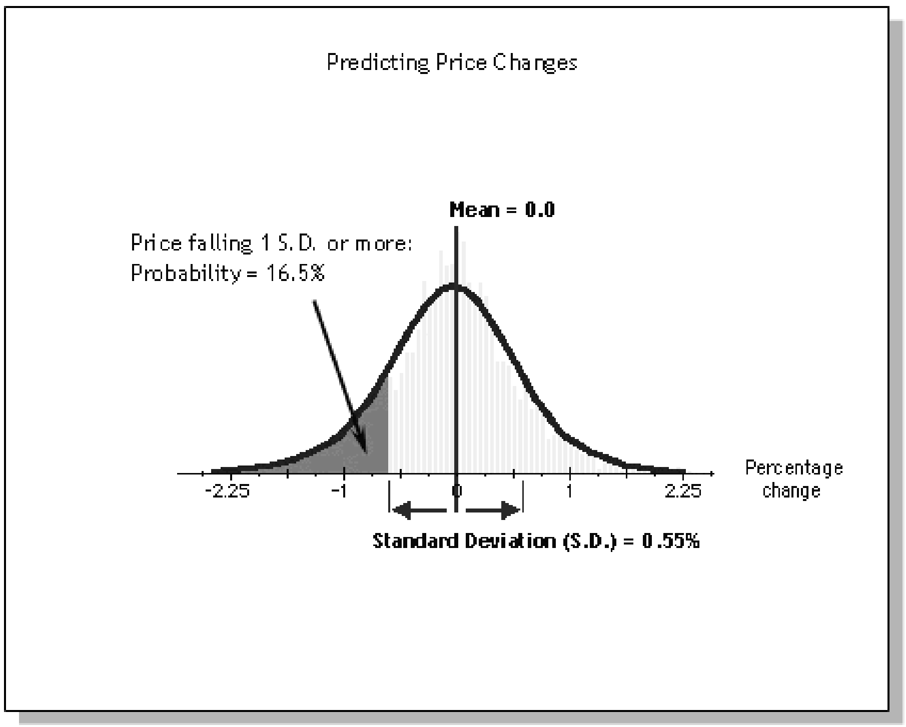

On daily VAR calculations, the mean return on assets is typically not significantly different from zero, so common practice is to fit a normal distribution around a mean return of zero.

The solid line in the above chart indicates the number of observations that we would expect to see in our sample if the price behaviour in this market was perfectly normal. The fit is reasonably close.

We can estimate the probability of the daily price falling by any given number of standard deviations by looking at the area under the fitted normal distribution curve, as shown in the figure, or we can read it from a standard normal probability distribution table which you find at the back of most statistics text-books.

[3] Short for autoregressive conditional heteroscedasticity!

The table below summarises some of the most popular probability levels used in financial risk management.

Probability	Price fall Larger than	In this case (μ = 0.0, σ = 0.55%)	Confidence Level
50.00%	μ	= 0.00%	50.00%
15.87%	(μ - 1.00 x σ)	= -0.55%	84.13%
4.95%	(μ - 1.65 x σ)	= -0.91%	95.05%
2.28%	(μ - 2.00 x σ)	= -1.08%	97.72%
0.14%	(μ - 3.00 x σ)	= -1.10%	99.86%

Thus, assuming this asset's return is normally distributed, the table shows for example that on any given day:

- The probability of its price falling by 1.10% or more (3 standard deviations) is only 0.14%
- Or (another way of saying the same thing) we can be 99.86% confident that that the market price on this asset will *not* fall by more than 1.10%

Daily VAR = Portfolio value x Z x σ

Where
σ = Standard deviation of return (i.e. volatility) in decimal
Z = An **adjustment factor**, representing the number of standard deviations (SDs) that corresponds to the required confidence level in a normal distribution (e.g. 1.65 SD for 95% confidence, 2.33 SD for 99% confidence, etc.)

Also known as: **Z score**

Example

Security: 10 year French government bond
Daily price volatility: 0.55%
Position: Long EUR 500 million

? What is the daily VAR on this position at the 95% confidence level?

Analysis
Assuming a normal price distribution, we can say with 95% confidence that on a single day the price will not fall by more than 1.65 standard deviations.

Daily VAR = 500,000,000 x 1.65 x 0.0055
= EUR **4,537,500**

In the language of risk management, we have approximately EUR 4.5 million of value at risk at the 95% confidence level.

But please note that VAR is a probabilistic estimate, not a certainty. Even then, remember that the estimate is based on the *assumption* that price movements in this market are normally distributed, which of course may not be the case in reality.

The idea behind VAR is that a trading house should make sufficient capital provision to cover potential losses on all but the most extreme scenarios. Covering losses at the 95% confidence level still means that we risk losing more than that amount, on average, one day every month! If this is too often for comfort, we may want to set aside more capital to cover VAR at a higher confidence level - say 99%, which is indeed that confidence level at which banks using their own internal risk models are required to calculate their **market VAR**, according to the Bank of International Settlement's 1996 Basle Amendment[4].

3.2. Diversified VAR

The simple VAR methodology can be easily adapted to allow for the portfolio effect that we saw in section *Passive Portfolios*: the reduction in specific risk that results from holding a diversified portfolio of more or less uncorrelated instruments. If the return on each security is normally distributed, then the return on the total portfolio will also be normally distributed. Therefore, we can derive the **diversified VAR** on a portfolio by first calculating the risk on the portfolio using the formula introduced in section *Return and Risk*.

Example

Suppose you are long EUR 500 million in the following mixed portfolio of French securities:

	A: 10-year bond		**B: CAC-40 equities**
Position	Long		Long
Market value	EUR 300 million		EUR 200 million
Daily price volatility (σ)	0.55%		1.12%
Correlation coefficient		+0.74	

? What is the daily VAR on this portfolio at the 95% confidence level?

Analysis

Step 1:
Calculate the risk on the portfolio, taking into account the correlation of returns between these securities:

Portfolio risk $= \sqrt{ (w_A \times \sigma_A)^2 + (w_B \times \sigma_B)^2 + (2 \times w_A \times w_B \times \sigma_A \times \sigma_B \times \rho_{AB}) }$

Where:
w_A = Proportion of the portfolio invested in asset A - the 'weight' of A
w_B = Proportion of the portfolio invested in asset B - the 'weight' of B
σ_A = Standard deviation ('risk') of the return on asset A (in percent)
σ_B = Standard deviation of the return on asset B
ρ_{AB} = Correlation coefficient of returns between assets A and B

Portfolio risk
$= \sqrt{ (3/5 \times 0.55)^2 + (2/5 \times 1.12)^2 + (2 \times 3/5 \times 2/5 \times 0.55 \times 1.12 \times 0.74) }$
= **0.73%**

[4] *Amendment to the Basle Accord to Incorporate Market Risk*, published in 1996 by the Bank for International Settlements (BIS), the Basle-based organisation whose members are the central banks of the main OECD countries.

Step 2:
Calculate daily VAR on the portfolio by slotting its standard deviation into the simple VAR formula that we developed earlier:

Daily VAR = Portfolio value x Z □x σ_P

= 500,000,000 x 1.65 x 0.0073
= **EUR 6,022,500**

Simple Vs. Diversified VAR

The diversified VAR on this portfolio at the 95% confidence level is EUR 6,022,500. This is lower than the sum of the **simple VARs** on each of the positions in the two securities taken separately:

- Long EUR 300 million in the 10-year bond
- Long EUR 200 million in the French equities

a) Enter the simple VAR at the 95% confidence level for each of these two positions taken separately in the boxes below (to the nearest EUR) to confirm this.

Daily VAR of the bond position alone = EUR

Daily VAR of the equity position alone = EUR

Sum of the two separate VARs = EUR

4. VAR and Time

? Once we know a market's daily VAR, how do we assess the market risk on a position that is held for longer time periods?

Example

Settlement date:	5 February
Security:	French CAC-40 index portfolio
Daily price volatility:	1.12%
Position:	Long EUR 500 million

Suppose it takes a minimum of 5 working days to unwind this position. What is the VAR over that period at the 99% confidence level (covering 2.33 standard deviations)?

Analysis
The VAR over 5 days will obviously be higher than the daily VAR, but probably less than 5 times that figure because it is unlikely that a 'down day' in this market will be followed by 4 more down days in succession.

Square root of time rule: if the return on an asset or portfolio is normally distributed, then:

VAR over T days = Daily VAR x $\sqrt{T}$

This rule can be easily derived from the formula for the standard deviation of the return on a portfolio of assets that we presented in section *Portfolio Return and Risk*, above, if we imagine that each day that the asset is held represents one element of a 'portfolio' of T uncorrelated assets. However, such a derivation lies outside the scope of the IFID Certificate syllabus. All you are required to do is be able to apply the rule and in this case:

5-day VAR = (500,000,000 x 2.33 x 0.0112) x √5
= **EUR 29,125,000**

The square root of time rule indicates that the market risk on a position increases less than proportionately with time, as illustrated in the figure below which shows the VAR band on the CAC-40 index looking forward from 5 February onwards.

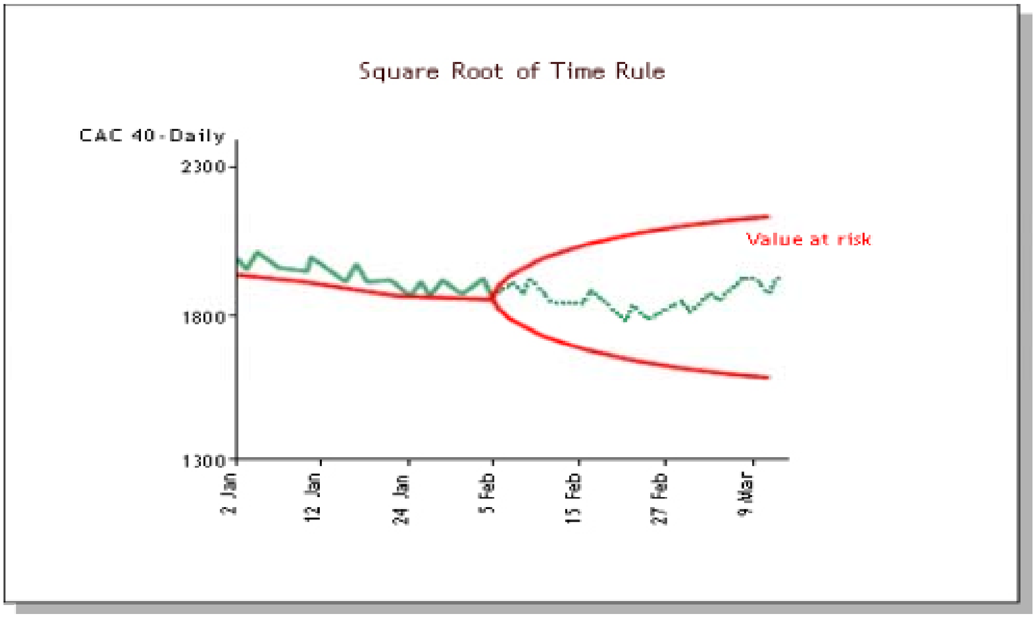

The same rule describes the relationship between an option price and its time to expiry, since option pricing models make the assumption that the return on the underlying instrument is normally distributed (see Option Pricing and Risks - Theta). But note:

Strictly speaking, the square root of time rule only applies in markets that display normal return distributions, which implies zero serial correlation.

The higher the degree of serial correlation in a market, the more proportional risk is to time and the more the cone in the figure above stretches out into two straight lines.

5. Limitations

Key Assumption of VAR

- The returns on financial instruments are assumed to be joint-normally distributed with constant standard deviation
- So we can use the estimated volatility of a trading portfolio to predict monetary losses with varying degrees of confidence
- And on that basis we allocate sufficient capital to cover trading losses in all but the most extreme market scenarios.

The critical assumption in all this is that the distribution of returns should be normal. The worrying thing is that it implies that the markets follow a random path: that one day's price action has no bearing whatsoever on the next day's price, or on any other day's.

This means that a monkey throwing darts at the Wall Street Journal is as likely to pick market winners as is Warren Buffet!

! Returns on financial instruments are assumed to be joint-normally distributed. This implies that the markets are random!

The Evidence

Most empirical tests of the actual behaviour of market prices have come up with some important results that limit the power of **parametric VAR calculations** (i.e. those based on the standard deviation of the market's return) such as the ones that we performed earlier in this section.

Fat tails and tall peaks

Price distributions in many markets tend to display **fat tails** (technically known as **leptokurtosis**). This means that there are more observations of extreme price changes (**outliers**) than would be predicted by the normal distribution. At the same time, the peaks around the mean of some price distributions appear taller than normal (**kurtosis**).

In other words, when the markets are not going crazy they seem to spend rather too long doing very little.

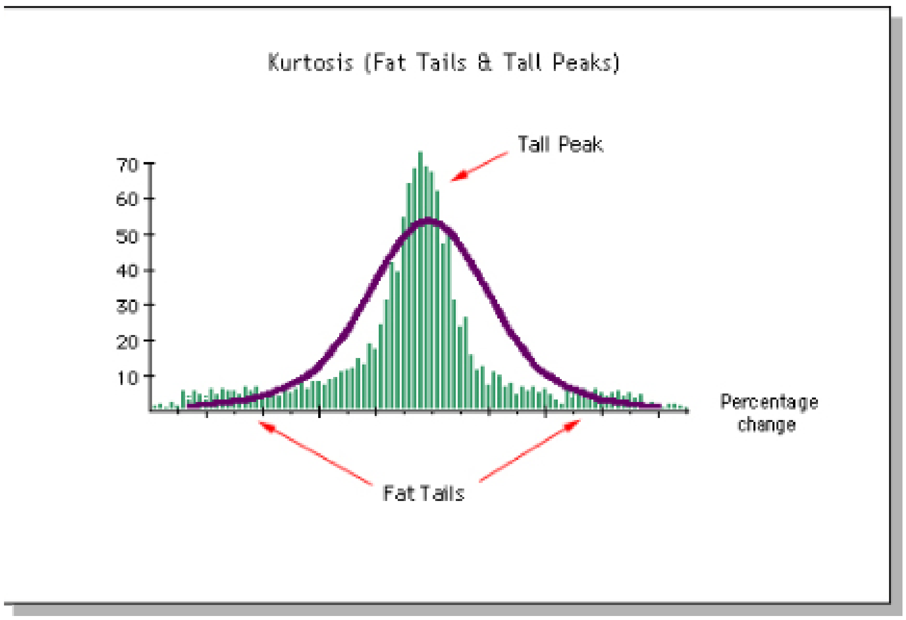

Fat tails indicate that bull markets have a tendency to gather momentum and, conversely, when things get bad they tend to get really bad. So VAR may underestimate the risk that is actually present in some markets. If the markets were really normally distributed, then the shocks caused by the emerging markets crisis of 1997/98 should only happen once every 5-7 thousand years; and a stock market crash such as the one of October 1987 only once every 3.5 million years!

The existence of fat tails also suggests that in extreme conditions the price action may no longer be normal and becomes instead serially correlated. This means that future prices are influenced by past history and therefore a down day could well be followed by one, or more, successive down days. In this case, as we said earlier, the square root of time rule would underestimate the longer term risk on a position.

Skewed distributions

The shape of the return histograms in some markets tends to be skewed, rather than symmetric like the one for the French government bond that we saw at the start of this section. In other words, in certain markets there are more recorded instances of down days than of up days, so any VAR calculation that assumes a normal distribution would not give an accurate assessment of the risks that are actually present.

Another notable example of skewed markets are the ones for credit spreads on corporate bonds, where the probabilities of default tend to be very low but the losses in such events tend to be very large indeed. This implies that the price distributions on products such as credit default swaps (which as we explained in module Credit Derivatives – CDS Pricing tend to be derived from yield spreads of bonds over swaps) do not follow normal distributions and any calculation of the VAR on a CDS position based on standard deviations and normal distribution assumptions would seriously underestimate its true risk.

A detailed discussion of the types of model used to calculate credit VAR lies outside the scope of the IFID syllabus.

Variable correlations

Price correlations between different markets are neither stable nor **linear**, as they should be if the markets were truly joint-normally distributed. For example, the left panel in the figure below shows the historic return on an asset against the return on its option. It is in the nature of options that their correlation with the underlying is not linear (see Option Pricing and Risks - Delta).

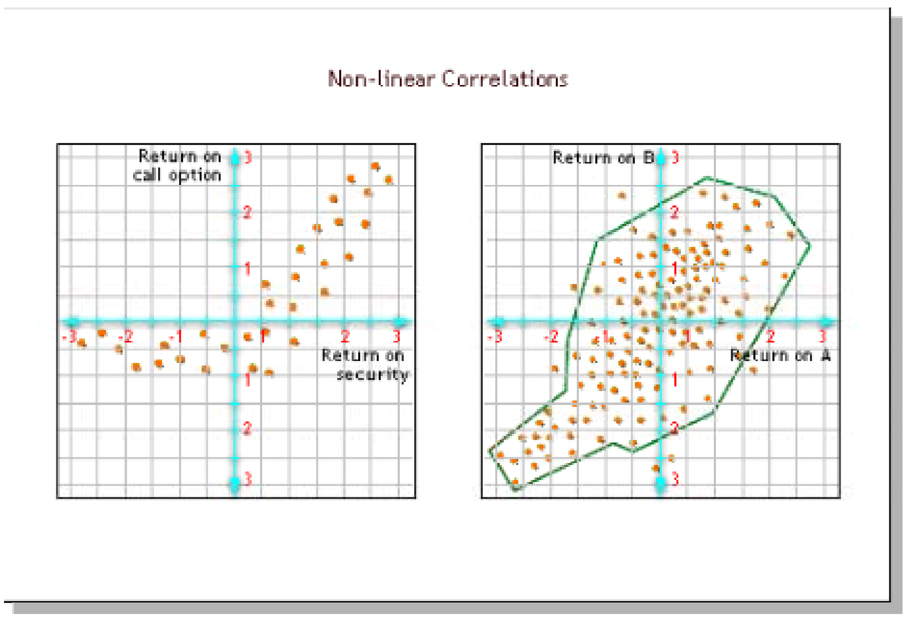

The returns on the two securities in the panel on the right appear to be reasonably uncorrelated in normal market conditions, but the appendix-like form at the bottom-left suggests that their correlation may in fact increase in a bear market, which is something that has often been observed.

Volatilities are themselves volatile

Finally, it is worth remembering that market volatilities themselves tend to change over time (a phenomenon that is reflected in the options markets) and therefore so do the calculated VARs.

One phenomenon which has been observed in many markets is the tendency of periods of high (or low volatility) to persist over time; a phenomenon referred to as **volatility clusters** or ARCH[5].

5.1. Conclusion

Value at risk can be a powerful tool of risk management, but it relies critically on the assumption that the returns in different markets are joint-normally distributed. In some markets, we know that these conditions may not hold, and that could make value at risk calculations potentially misleading. So where does this leave us?

It is tempting to look for more complex statistical distributions with which to model the real markets. But researchers have found this route hard-going computationally, with little or no gain in predictive power. This is why VAR, with all its limitations, remains the foundation of most risk management systems today.

In practice, the VAR calculations must be supplemented with stress tests of positions under selected market scenarios, where the risk manager will simulate the likely trading losses in each case.

In this section, we focussed on the basic techniques of measuring and predicting risk. At the end of the day, the practice of risk management remains very much an art. In practice, parametric VAR calculations must be supplemented with stress tests of positions under selected market scenarios.

[5] Market scenarios for the performance of stress testing may be selected in different ways:

- **Historic simulations**
  Calculating the profit/loss that would have resulted if a current trading position was marked to market at the rates prevailing at some notable date in the past - e.g. Black Monday, 1987. The idea here is to assess what would be the effects of a repeat of the same market conditions at some point in the future.

- **Selected scenarios**
  Calculating the profit/loss that would result on the current position assuming a specific set of market prices. This technique allows the risk manager to assess the implication for the firm of some imaginable 'nightmare scenario'.

- **Monte Carlo simulations**
  Simulating the future paths of asset prices using a random number generator and calculating the profit/loss on a position at the end of each simulated price path. The behaviour of the simulated path is a function of the assumed volatility of the asset price.

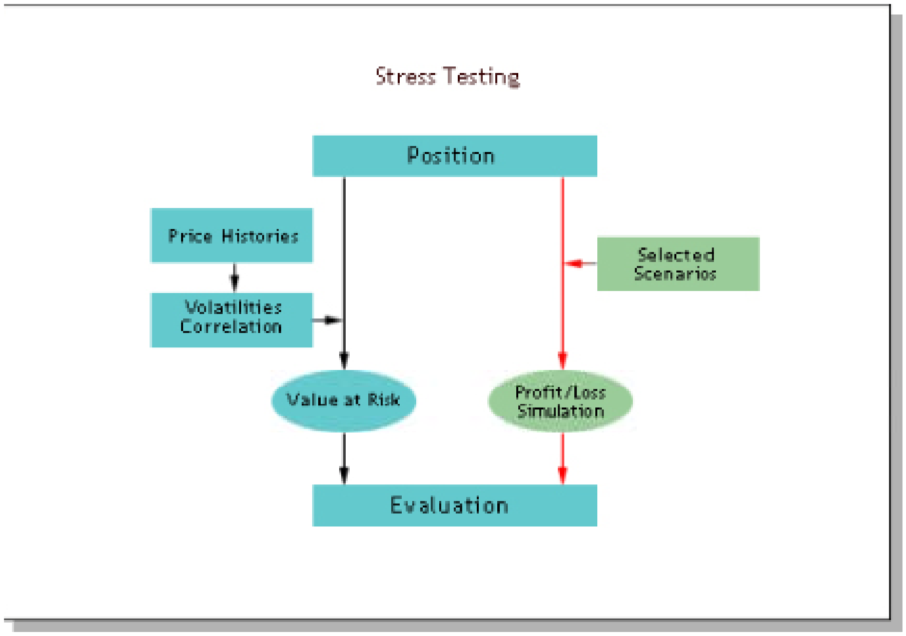
Stress Testing
Position
Price Histories
Volatilities
Correlation
Selected
Scenarios
Value at Risk
Profit/Loss
Simulation
Evaluation

BIS VAR rules
Summary of the rules allowed by the Bank for International Settlement concerning the use of internal VAR models.

In 1996 the BIS published its *Amendment to the Basle Accord to Incorporate Market Risk*. This was motivated by the increased proprietary trading activities of banks under its supervision and for the first time allowed banks to use their own internal models to calculate their daily VAR.

Briefly, the *Basle Amendment* states the following:

- Position VARs should be calculated daily using a 99% confidence level

- The VAR models may take into account the return correlations both within and across different asset classes (i.e. the extent to which interest rate positions might offset the market risk on, say, FX positions) provided the national supervisory authorities are satisfied that the institution's system for measuring correlations is "sound and implemented with integrity"

- A minimum of one year's historic market data must be used to estimate the daily volatility and correlation of returns in the various asset classes

- All trading positions are assumed to be held for a minimum period of 10 days and their 10-day VARs can be extrapolated from their daily VARs using the square root of time rule

The amount of capital that a bank needs to have in order to cover its market risk will be the greater of:

- **The previous day's calculated VAR at the 99% confidence level; or**
- **3 times the average VAR calculated over the preceding 60 business days**

Banks are also required to maintain qualitative standards, as follows:

- There should be an independent Risk Control Unit within the bank that reviews the risk measurement and management systems at regular intervals

- The VAR calculations should be supplemented with routine and regular stress-tests on trading positions, although no specific guidelines were given as to how exactly this should be done

The *Amendment* recognises that banks using their own internal VAR models have an incentive to understate their capital requirements, so the internal models used are subject to periodic **backtesting**. This involves comparing the financial losses predicted by the models over a period of time with the actual losses incurred, and banks that are found to have underestimated their capital requirements are subjected to punitive charges by their national supervisors, depending on their margin of error.

6. Exercise

6.1. Question 1

In the language of risk management, we have approximately EUR 4.5 million of value at risk at the 95% confidence level.

a) Which of the following is the correct interpretation of VAR in this case?

- ☐ We will lose more than EUR 4.5m on 5 days out of every 100
- ☐ We will lose less than EUR 4.5m on 95 days out of every 100
- ☐ We will never lose more than EUR 4.5m on any day
- ☐ We can be 95% confident that the trading loss in one day will not exceed EUR 4.5m

6.2. Question 2

The implied volatility of the options on the FTSE-100 index of UK equities is currently 18.0% per annum.

a) Based on this information, calculate the daily VAR on a GBP 55 million portfolio of UK equities:

- At the 99% confidence level (i.e. 2 standard deviations)
- Assuming the year has 262 working days

Enter your answer in the box below, in sterling to the nearest pound, and validate.

6.3. Question 3

A US investor has a portfolio of French bonds and equities worth EUR 500 million, whose currency risk has not been hedged. The investor's position and the characteristics of each market are summarised below (in all cases we assume the daily mean return on these markets is zero):

		A 10-year Bond	B CAC-40 Equities	C EUR/USD FX
Position		Long	Long	Long
Market value (EUR millions)		300	200	500
Daily % price volatility (σ)		0.55	1.12	0.65
Correlation Matrix	**A**	1.00	+0.74	-0.49
	B	+0.74	1.00	-0.52
	C	-0.49	-0.52	1.00

Note: The correlation matrix shows the correlation coefficients between all pairs of asset classes. For example:

- Correlation of returns between bonds and equities is +0.74: there is quite a tendency for both markets to rally simultaneously
- Correlation of returns between bonds and FX is -0.49: there is a tendency for a rallying bond market (lower interest rates) to be associated with a weakening EUR (stronger USD), but the association is weak
- Correlation of returns between equities and FX is -0.52: there is a stronger tendency for a rally in French equities to be associated with a weaker EUR

If the EUR/USD spot rate is 1.0490, what is the daily VAR on this portfolio, in USD, at the 95% confidence level (i.e. 1.65 standard deviations)?

a) First, calculate the risk on the portfolio using the formula developed in Portfolio Construction - Portfolio Risk & Return.

Portfolio risk $= \{ \Sigma (w_i \times \sigma_i)^2 + \Sigma\Sigma (2 \times w_i \times w_j \times \sigma_i \times \sigma_j \times \rho_{ij}) \}$

for $i = 1 \ldots 3$, $j = 1 \ldots 3$, $i \neq j$

Where:
w_i = Proportion of the asset portfolio value committed to instrument i
σ_i = Standard deviation of return on asset i (in percent)
ρij = Correlation coefficient of returns between assets i and j

Note that here the portfolio weights are, respectively, 3/5, 2/5 and 5/5, as the FX exposure is on 100% of the underlying bonds and equities.

Enter you answer in percent to 2 decimal places.

b) Next, apply this portfolio volatility to the VAR formula developed in section *Simple VAR* (assuming a mean daily return of 0).

VAR = Portfolio value x Z x σ_P

Where:
σ_P = Portfolio risk (standard deviation of portfolio return, in decimal)
Z = The number of standard deviations corresponding to the required confidence level (e.g 1.65 for 95% confidence)

Calculate the VAR on this portfolio in EUR, to the nearest Euro.

c) Finally, express the VAR on this portfolio in USD equivalent to the nearest dollar, at EUR/USD 1.0490.