4.3.1 Risk Evaluation

Unit: Unit 1 – Treasury Strategy

Module: Module 4 – Financial risk

Date: 1 September 2016

Summary: Many techniques are available for the evaluation of risk, but effort needs to be made to ensure that a suitable measure is adopted and that the scale of the risk is most important

Keywords: evaluation, key risk indicator, key control indicator, value at risk, enterprise risk management, feedback, cash flow at risk, sensitivity analysis, scenario analysis, stress testing, reverse stress testing, randomness, uncertainty, correlation, portfolio, mean variance, Monte Carlo, stochastic, tail risk, volatility,

Contents

1 Introduction 1
2 Why do we evaluate risks? 1
2.1 Key risk indicators (KRIs) and key control indicators (KCIs) 3
3 What measure for evaluation? 3
3.1 Alternative measures of risk 4
3.2 Dimensions of risk measurement 5
3.2 Impact-only approaches 8
4 A history of risk evaluation 9
5 More on impact-only methods of evaluation 10
5.1 Sensitivity analysis 10
5.2 Scenario analysis 11
5.3 Stress testing 12
5.4 Reverse stress testing 12
6 Randomness and uncertainty 13
6.1 Statistics and markets 13
6.2 Uncertainty 13
7 Statistics 14
8 Correlation and portfolios 14
9 Mean variance model 15
10 Monte Carlo and stochastic methods 17
10.1 Stochastic modelling 18
11 Presenting the data 19
12 What can go wrong? 20
12.1 Correlation measurement errors and changes 21
12.2 Risk management in silos (enterprise risk management failure) 21
12.3 Different measures of risk adopted 22
12.4 Tail risk 22
13 Trusting the market 23
14 What you can hope for in evaluation 25

© Association of Corporate Treasurers

Learning outcomes:

1. Evaluation poses question of risks size compare to the firm and each other, ranking risks
2. Evaluation allows risks to be taken as well as controlled
3. The concepts of tolerance, appetite and budgeting depend on evaluation
4. Key Risk Indicators (KRIs) and Key Control Indicators (KCIs) are parts of empowerment and control
5. There are many possible measures for risk, from eps and shareholder value to money amounts and credit risk measures
6. Measuring investments in terms of shareholder value (NPV) invites risk to be also measured in that way, perhaps in an ideal world
7. Probability and Impact remain popular evaluation tools – sometimes they can be combined to make one measure of risk, and sometimes not. One measure of risk could be seen an indication of successful evaluation
8. Measurement techniques include sensitivity analysis, scenario analysis, stress testing and reverse stress testing
9. An understanding of statistics is useful for risk management
10. The benefits of diversity underpin the ideas of Enterprise Risk Management and many modern techniques for risk evaluation. The strengths and weakness of this approach are crucial to its understanding
11. Mean variance models are useful desktop tools for a treasurer
12. More modern and flexible approaches to modelling include Monte Carlo techniques
13. Careful thought needs to be given to presentation of evaluation results
14. Evaluation is not without its own risks, including issues over correlations between risks, risk evaluation being performed in silos, different measures of risk being adopted and tail risk
15. The behaviour of financial markets is important for treasurers to understand because of their heavy use by treasurers, how they can be wrong but how they also carry information
16. Ranges of outcomes are a useful approach to planning and risk measurement generally

© Association of Corporate Treasurers

1 Introduction

Evaluation is where risk management starts to get difficult. It aims to answer two important questions:

- How big is each risk in comparison to the firm (and hence its ability to cope)?
- How big is each risk in comparison to each other risk?

Answering these two questions will get us straight to dealing with important issues such as risk appetite and budgeting.

While these are straightforward questions to pose, they immediately throw up further challenges, the main ones being:

- What suitable measures are there?
- Can any one risk be compared with every other risk? Is it possible to have a ranking of risks?

This reading will takes us through many of the alternative measures available and show some of the techniques that can reach those measures. Perhaps some risks will not respond as well to some measures as others, but it must be borne in mind that the aim is to achieve an overall ranking. However, if it is not possible to do this, then some way has to be found to deal with this, perhaps with a different risk appetite for different types of risk. In MCT we are mainly thinking about financial risk evaluation but treasurers must bear in mind that business risks are on the agenda too.

There are three complementary appendices for this reading (Statistics, Portfolio theory and Value at Risk). The appendices cover much of the detail that is discussed strategically in this reading. Treasurers need to be familiar with the language and techniques.

2 Why do we evaluate risks?

At the evaluation stage in the framework we quantify the risk as far as possible, in order to design a response to that risk. Without a measure of a risk, we cannot decide on the importance of a response to it, and if a response is chosen that aims to reduce that risk in some way, we need to know that the response is successful in reducing the risk. And if we chose simply to monitor a risk, then it needs continuous measurement to check that it has not grown beyond what is acceptable.

This also works in reverse. If there are risks that we actively seek in a firm, to raise the return, then we may need to take more risk. For example, if an important product becomes commoditised, so that several different manufacturers supply it and the margin falls, then that element of the business has fallen in both risk and return. It may be time to invest and produce a modified and better product which will require the taking of risk (e.g. capex and marketing spend) to achieve more reward.

© Association of Corporate Treasurers

Therefore, to accommodate the principles of risk tolerance, appetite and budgeting, we must create a common measure between risks. It could be argued that because some risks are so different in nature, it is not practicable to apply a common measure to all risks. However, Enterprise Risk Management (ERM) then becomes more difficult. It is possible, however, to have multiple risk appetites, each corresponding to different types of risk.

It is useful to consider here how banks and insurers approach (and indeed are required to approach) the management of risk. As a broad summary, banks are required to allocate capital to every risk that they have. The Basel II regulations identify, for banks, three different types of risk (liquidity risk is considered by Basel III with its approach to liquidity management):

- Credit risk
- Market risk
- Operational risk

With credit risk, for example, each asset is risk weighted and a specific amount of capital is required to support that loan.

Example 1: Allocation of capital

Bank 365 makes loans to Goody 2 Shoes, an A (investment grade) rated corporate and to Last Chance Saloon, a B (speculative grade) rated corporate. Under Basle II regulations, capital is required to be allocated at 8% of risk weighted assets[1]. The risk weightings set out by management and regulators is as follows:

A rated	25%
B rated	150%

For each $100 of loan made to these borrowers the capital required is calculated:

Goody 2 Shoes, A rated	100 X 0.25 X 0.08	=	$2
Last Chance Saloon, B rated	100 X 1.50 X 0.08	=	$12

Note that this example is simplified and excludes the effect of Loss Given Default, another major factor.

We can see that all the elements of risk tolerance, appetite and budgeting are present in this approach. Risk tolerance could be limited by the total amount of capital available, risk appetite is how much capital the management actually want to use and budgeting is quite easy to do, by simply allocating a strict capital amount to each department, X to Europe credit, Y to US credit, Z to market risk in treasury and so on.

One significant difference present for banks and insurers, however, is that they both have large portfolios. The allocation described only makes real sense if Bank 365 is able to make many loans to borrowers. It can then afford, in extremis, for 2% of A rated borrowers to fail and 12% of B rated borrowers to fail. This should be approximately how many borrowers do

[1] Capital allocation under Basle III is rather more complex.

fail, based on experience. What is happening here is that as long as losses in the long run accord with experience, then the loss rate can be priced into the operating business model and there is essentially no uncertainty, no risk from losses in themselves (there may be an issue with competitiveness of course). Similarly, insurers with, say, car insurance, will collect sufficient premium to pay for claims, plus some extra for administration and profit. There is little risk, despite the business operating in risk management.

For corporates with foreign exchange risk, for example, there may be no large portfolio in which case any risk is proportionally larger, implying that corporations are not naturals at absorbing foreign exchange risk. They may, however, take on risk in production, or a core skill, where there is volume.

2.1 Key risk indicators (KRIs) and key control indicators (KCIs)

Evaluation is also required to use other important elements of the framework, namely:

- Key risk indicators (KRIs)
- Key control indicators (KCIs)

These reflect essentially the two major components of risk appetite. They are useful tools when evaluating risks, especially in the context of change and in responding to risks, or indeed empowering staff to take risks.

Example 2: KRIs and KCIs

Asimov Inc operates in the capital goods sector and regularly makes tenders in foreign markets. The tenders are based on foreign exchange rates prevailing at the time of bid but there is risk in the tenders which is difficult to hedge because it is not certain whether a particular tender will be accepted.

The *Key Risk Indicators* for Asimov might comprise:

- Number of bids outstanding
- Size of exposure in bids outstanding
- Length of tender periods
- Volatility of currency pairs involved in tenders

The *Key Control Indicator* for Asimov might comprise:

- Maximum exposure to any one currency pair, adjusted for volatility (and hence time)

Thus Asimov makes a series of measures on the bids outstanding, recognising that risk increases as KRIs increase. There is no desire to limit the ability of the business to take on work or submit tenders and so control is exercised by limiting the maximum exposure, and taking appropriate hedges in the financial markets to keep the total risk within limits.

3 What measure for evaluation?

It has to be said from the outset that measuring risk is not easy. However, treasurers (and indeed managers) are not appointed to do easy things. Taking the different risks that face a firm and trying to find a common measure is always going to be a challenge. Putting reputational risk or legal risk alongside foreign exchange or interest rate risk and expecting a measure to fall out is not simple. It might actually be asking too much to find a common measure and many firms will use a mix of measures for different risks. However, to properly evaluate risks with a view to establishing how much risk a firm wants, a common measure is exactly what is needed.

If the object of a firm is to improve shareholder wealth, then the choice is easy. Risk is measured by its impact on shareholder value. This should be the gold standard and deviation from this should only be with good cause. However, this can be difficult to conceive when considering offering discounts to regular customers or considering the pros and cons of a day to day decision in treasury, such as rollover periods. It is nevertheless a good backdrop. After all, we know that the measurement of investment decisions is all around finding a Net Present Value, it therefore follows that risk measurement should be about variation to Net Present Value.

Shareholder value is quite easy for trained finance professionals to understand and even implement as a universal system, but other managers may not find it so easy.

3.1 Alternative measures of risk

The table in Exhibit 1 (which we have already seen) offers some alternatives for the measurement of the impact for a particular risk.

Exhibit 1: Measure of risk impact

Measure	Impact of the event
EPS	On earnings per share for current or future periods
Shareholder value	Straight impact on share price through discounting all future cash flow lost
Cash flow at risk (see Exhibit 2)	Looks explicitly at cash flow rather than valuation - for current or future periods
Sales impact	Loss in sales
Margin	Reduction in margin
Working capital	Impact on working capital
Financial covenants, e.g. interest cover or Debt / EBITDA	Impact on covenants and proximity to breach
Money amount	A simple monetary amount (e.g. as indicated by VaR approach)

Each of these might be more appropriate at particular times. If a firm is under borrowing stress, then financial covenants might be the measure exercising management for every waking hour. Any decision might be based around that. For example, a firm running close to a borrowing limit (limited by facility size or covenant) will prioritise managing the cash, knowing that the risk is possibly insolvency.

For managers under pressure from investors, EPS might be important, although this might also be managed by 'sales impact' or by 'margin impact'.

Exhibit 2: Cash flow at risk

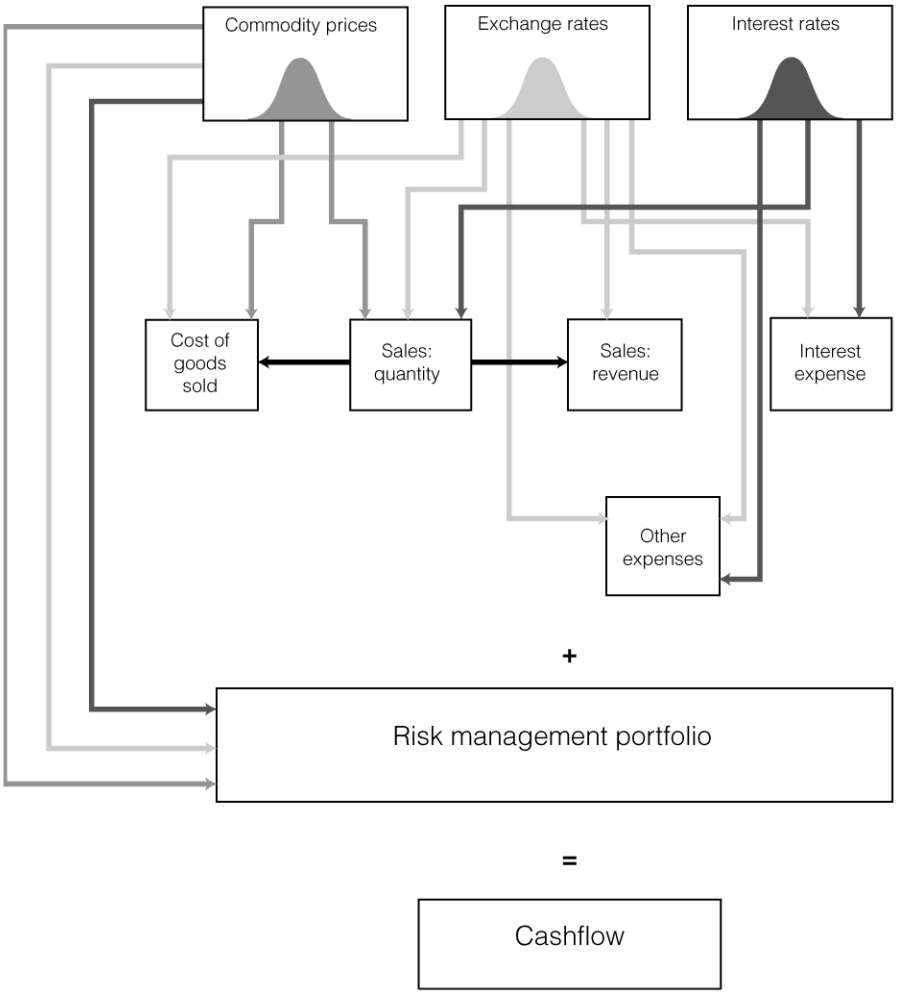

3.2 Dimensions of risk measurement

The above measures are suitable components of the impact part of the risk map. Several of the measures, however, could also be considered as an overall measure of risk, combining probability and impact to give an overall measure against which all risks can be compared. Value at Risk, in particular, combines probability with impact to give a specific size to a risk.

Value at risk (VaR) is dealt with in detail in Appendix C. In the context of this reading however, we will note that VaR is a statistical technique which summarises for a financial risk:

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

VaR and associated techniques such as Cash Flow at Risk (CFaR) are perfectly aligned with a risk map, as they quantify both a probability and an impact for many risks. The VaR 'value' number

is then the answer to "How much risk am I living with for the period until I review my exposure again?"

VaR has a number of disadvantages for quantifying risk in non-financial corporations, so using VaR is not always meaningful. Outside VaR techniques, if we classically seek a two dimensional measure of risk to allow a mapping of that risk with others, then a simplified measure of that risk might be impact multiplied by probability. However even this simplified approach is not always practical or helpful and might only be suitable in certain situations.

Example 3: Plotting the effects of market price movements on the risk map

Kinesis plc (functional currency GBP) has made a one off sale to the US and is owed USD 6 million. Until the USD are sold Kinesis has a transaction exposure in that amount. The current exchange rate is 1.50. The contribution from the sale at the current exchange rate is GBP 400,000. Kinesis is looking for a way to measure that risk so that the risk management framework can be followed. However, the risk will not easily fit onto a risk map because the probability of an adverse event is not known. Beyond that, the impact increases the further the exchange rate moves from 1.50 and therefore the risk could either be plotted as:

- High probability (of a small movement) and hence low impact, or
- Low probability (of a high movement) and hence high impact.

In fact, the complete risk is best plotted as a line across the risk map.

Kinesis decide that the minimum contribution from the sale to meet required shareholder returns is GBP 100,000 so they wish to deal with it on that basis. The minimum GBP100,000 would require a move in the exchange rate to 1.6216.

Kinesis have no idea of the probability of this event, although their bankers tell them that the foreign exchange options market prices a move to 1.6216 at 0.15, or 15% chance within the next three months. This allows Kinesis to plot the risk on the risk map - the 'event' is that GBPUSD moves greater than 1.6216; the probability is 15% over the next three months, and the impact is a £300,000 loss of contribution.

Kinesis can now devise a risk response.

Example 4: Shortcomings in expected probability and impact

Yummy Foods Inc manufactures food products for Piggy Inc, a major retailing chain and is heavily exposed to their credit risk, as goods shipped to Piggy Inc are not paid for until 90 days after delivery. Piggy Inc contributes most of the turnover of Yummy Foods. The average size of the receivable is USD 50 million, and the probability of failure is assessed at 1 in 100 over a time scale of one year. If Piggy were to fail, then Yummy Foods would suffer heavy losses: the impact would be a cash loss of probably the whole of the receivable of USD 50 million, an extremely large sum for Yummy Foods, who typically run on small margins. The expected loss, representing a single overall measure of this risk, is 0.01 x $50 million = USD 0.5 million in a year.

Simplistically one could plot an expected risk of USD0.5 million at a probability of 100%, but this is not helpful here because Piggy's failure will either happen or not happen, rather than

partly happen, and Piggy is a major part of Yummy Foods' business. A better measure is to plot the 1% probability of losing the full $50 million and to manage the risk on this basis.

Example 5: Portfolio effect
Argonaut Inc sells food products to several hundred differently owned restaurants around the country, trading on open account. The total receivables book is about USD 20 million (about 3 months' sales) and no one exposure exceeds USD 75,000. Historical analysis shows that about 0.5% of restaurants fail in a year, causing average annual losses of about USD 85,000. This expected loss represents a single overall measure of this risk (expected impact $85,000 and probability 100%) and is used in the management of credit risk. Argonaut can manage this on a portfolio basis.

Example 6: Liquidity risk
Argonaut Inc, which is publicly owned, has relatively stable earnings and is considering a large dividend financed by leveraging up its balance sheet. It is advised that it could take its leverage up to a debt level of three times EBITDA, which is where its bankers say that senior lenders will currently be happy to lend. This ratio effectively becomes a measure of liquidity risk, because any higher leverage would make refinancing so difficult that the survival of the firm would be under threat.

Argonaut chooses to target a ratio of 2.5, so that some underperformance on earnings is not too dangerous, which also allows for the fashion in the banking markets to tighten slightly in the future. This ratio has become the banks measure of risk and in its turn, has become a key risk measure for Argonaut. It may be the case that this ratio becomes a financial covenant in a loan agreement, when the risk becomes tangible.

The strategy to lever up means that Argonaut may have to change its risk maps from a direct value measure such as EPS or shareholder value, to the more indirect debt : EBITDA measure.

Example 7: VaR as one overall risk measure
Past Glory plc operates a pension scheme for employees and carries the investment risk of the assets in the scheme. The value of the scheme is subject to many risks, such as on interest rates, equity returns, inflation, and longevity as well as the credit risk of the sponsor. Some of the risks are known to interact with each other and so the effect of no one risk can be isolated. To measure risk in the pension scheme Past Glory uses one overall measure, derived using Value at Risk, and takes action to reduce this number to an acceptable limit.

The two dimensional probability / impact approach can be summarised as follows in Exhibit 3.

Exhibit 3: Dimensions of risk

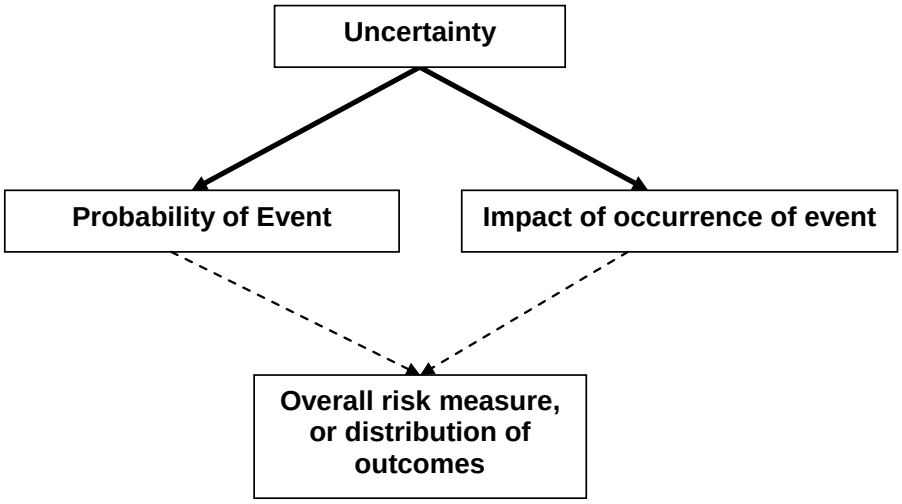

The examples show that some risks are best viewed in the two dimensions separately, so that risk management can be based on either managing probability or managing impact. Equally, one overall measure of risk may be better for example when risks interact and the probabilities of different risks are connected to each other.

3.2 Impact-only approaches

While an approach to risk measurement that includes probability is an ideal approach, sometimes probability is hard to define or needs to be layered in separately (e.g. by using Monte Carlo simulation). There are various 'impact-only' approaches and treasurers should understand how these might work. These include:

Sensitivity analysis: Or analytics, measures the change in value of an exposure for a given change in one variable, for example, market prices. Analytics are formula-driven, and give no indication of probability, but only of degree. Analytics can be used to show, for example, that a given bond portfolio worth $1 million now, will gain or lose value of $10,000 for every 10 basis point move in 5-year interest rates.

Scenario analysis: Scenarios attempt to indicate what might happen under different circumstances by answering a series of "what if" questions, i.e. by altering a number of variables. Scenario analysis should attempt to answer all possible "what if" questions, from the very probable but low impact events, to the unlikely but disastrous events.

Stress testing: This is a more extreme form of both sensitivity and scenario analysis (and so arguably not a new approach, but risk language includes the term) where more extreme events are tested against the relevant impact measurement adopted. This is common in the banking sector and can be applied against many measures. For example, in the banking sector, stress testing will be done for market risk, credit risk, operational risk and liquidity risk, and the stress tests are absolute benchmarks with no probability assigned at all.

Reverse stress testing: Moving up the catastrophe ladder, this sort of testing is to investigate what will break the firm. Again, it derives from the banking sector, where it looks at what might bring down the bank, but naturally it can be extended into non-financial corporates.

4 A history of risk evaluation

It is useful to see how risk evaluation has changed over the years and the two examples below are from the banking sector.

Exhibit 4: Financial risk measurement, market exposure

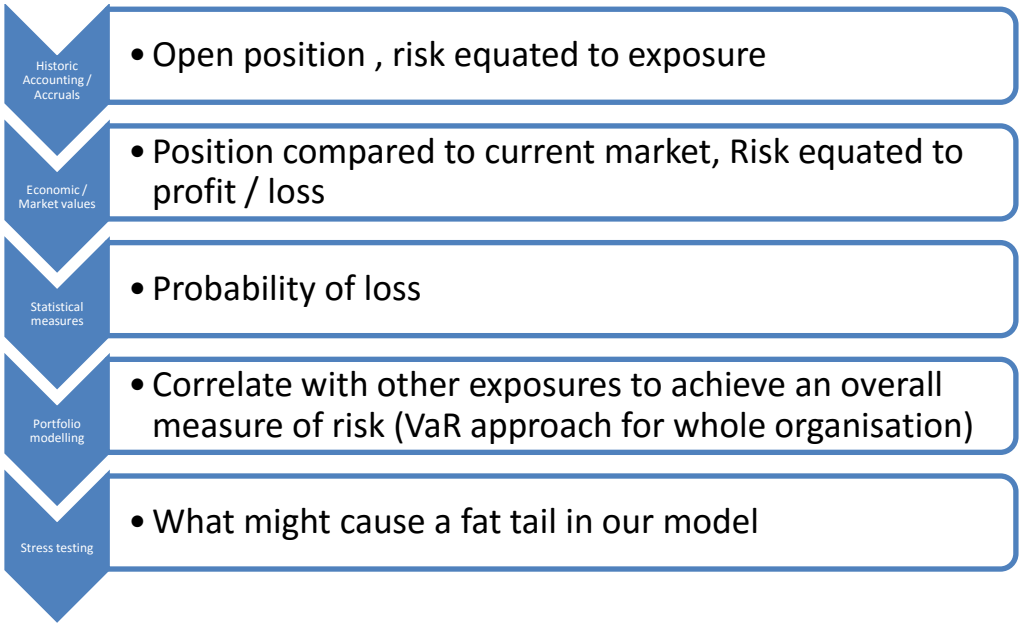

Exhibit 5: Financial risk measurement, credit risk in a lending book

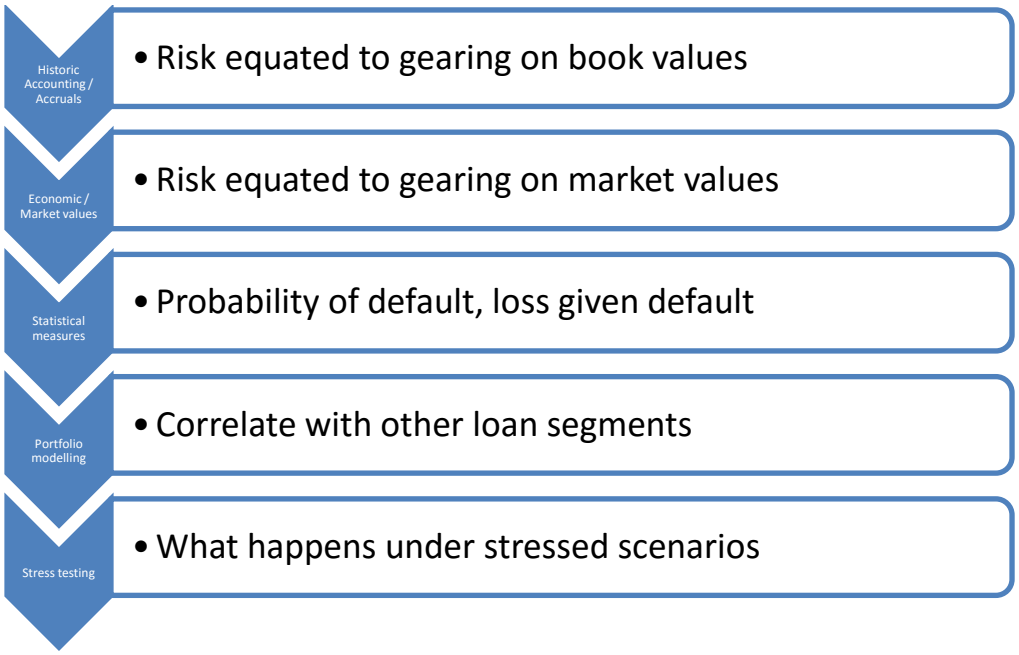

The general trend here is that simple measures have given way to sophisticated statistical techniques, only in their turn to give way to a more pragmatic approach which is arguably above the numbers to take an overall view of 'what could happen'.

It is easy to come to the conclusion that because statistical techniques (such as Value at Risk) were largely discredited in the credit crisis, then they are not worth considering in the corporate sector. However, to do so would be to ignore some extremely useful mathematics and indeed in advising the corporate sector, banks continue to use such techniques to show off their products.

5 More on impact-only methods of evaluation

This section looks at some evaluation techniques which are very useful but perhaps don't qualify as real measurements of risk in that they aren't really designed to contribute to the ideas of risk tolerance, appetite or budgeting. They are more indicators of what might happen, rather than a scientific measurement approach.

5.1 Sensitivity analysis

Sensitivity analysis involves changing one key input at a time and evaluating the effect of the change on the exposure. A good example is the "price value of a basis point" or PV01 when evaluating the sensitivity of bond prices (or indeed portfolios such as in pension schemes) to interest rate movements. PV01 is the amount by which the value of the bond changes, for a one basis point (i.e. 0.01%) change in interest rates. Sensitivity analysis can also be used in evaluating foreign exchange exposures – for instance quantifying the change in operating profit which would result from one foreign exchange rate changing by, say, 10% - a plausible move in the foreign exchange markets.

A two dimensional sensitivity analysis can also be quite helpful to see the effect of a combination of events. Three dimensional tables are also used and depicted in graphical form.

Example 8: Two way sensitivity analysis table

A UK based company owns a bond which is denominated in US Dollars. Current interest rates are 5.00% and the exchange rate between GBP and USD is 1.80. The GBP equivalent value of the bond is now £57.99 for every USD 100 face value of bond. The value of the bond according to different interest and exchange rates is as follows:

	4.00%	4.50%	5.00%	5.50%	6.00%
1.70	64.10	62.73	61.40	60.09	58.82
1.75	62.27	60.94	59.64	58.38	57.14
1.80	60.54	59.25	57.99	56.75	55.56
1.85	58.91	57.65	56.42	55.22	54.05
1.90	57.36	56.13	54.93	53.77	52.63

Data is for a five year, 6% coupon bond with semi annual payments.

On its own, sensitivity analysis quantifies an effect if an event happens, but doesn't provide an indication of probability. Outputs from sensitivity analysis can be fed into statistical

probability analyses to find the likelihood of the event actually occurring. Sensitivity analysis is however useful because users can put a magnitude on the risk against variables that they can understand. Sensitivity can be applied to very complicated models. For example, the profitability of a plant can be investigated according to sales volumes and, as long as the model is accurate, it can easily be discovered at what level the factory fails to be profitable or fails to cover overheads and so on. Whole company sensitivity to interest or exchange rates can also be easily calculated.

Sensitivity analysis can of course be carried to many dimensions, perhaps combined with statistical analysis either by assigning probabilities of scenarios happening, or by running computer simulations where individual inputs are varied – either independently or with some form of interdependence. Computer simulations can iterate the model thousands of times, varying inputs each time, and the results of such simulation techniques are then used to populate statistical analysis for a forward-looking basis. For simulations, means and standard deviations are not necessarily based on historical observation, but on the analyst's assumptions for future probabilities. Obviously these assumptions can be varied and interdependencies modified, so that a variety of models (typically at least a base case, a best case and a worst case) can be run and the results analysed.

Even simple combinations can have complex effects, and sensitivity analysis requires care if it is to be performed meaningfully.

5.2 Scenario analysis

Scenario Analysis is a different version of sensitivity analysis. Sensitivity analysis starts to look at "What If?" type questions but in a small and disciplined way, in line with how the relevant market might move. Scenario analysis takes this much further and asks more extreme questions to see what might happen to the firm.

Such questions might include:

- What happens if our bank goes bust?
- What happens if our treasury staff can't get to work?
- What happens if all the staff are sick?
- What happens if the share price collapses?
- What happens if the oil price jumps 50%, either up or down?
- What happens if our competitors bring out a new product?
- What happens if country X stops foreign exchange remittances?
- What happens if inflation jumps in country Y?
- What happens if the IT system breaks down?

The list of course can go on and on, but the origin of the questions is based on the risk assessment and mapping of risks completed in the assessment phase of the Risk Management Framework.

Scenario analysis, being similar to sensitivity analysis, suffers from the same weaknesses as sensitivity analysis, as no probability can really be assigned to these events, but it is

nevertheless useful, as these events do happen and one day will happen to some firm. Preparation will then have been key to the response.

5.3 Stress testing

Stress testing is really a more advanced form of scenario analysis and has reached prominence in the financial sector where the stresses are often proscribed by regulators when reviewing the ability of a bank to withstand certain shocks. Typical stresses might include:

- What happens if interest rates rise 1% in a parallel yield curve shift?
- What happens if the sovereign domicile is downgraded x notches?
- What happens if the base currency falls x%?
- What happens if depositors withdraw all 'on demand' funds?

In a corporate environment, the questions will be similar to those for scenario analysis.

5.4 Reverse stress testing

Reverse stress testing is simple in concept, and merely requires the identification of those events which would cause the business model to fail. It can easily be adapted to the corporate environment.

The point of reverse stress testing is that failure of the business occurs when customers and suppliers are unwilling to deal with the business – i.e. before insolvency. A key issue is how long it is between failure of the business and actual insolvency – and what can be done in this time?

The idea of reverse stress testing is to have a disciplined procedure to find weaknesses in the business, and hence what should be done to address them. Bear in mind that low probability 'tail risks' actually happen far more often than statistics predict. Identifying and planning for these extreme events gives you a better chance of dealing with them should they occur. Such events should be plausible scenarios, outside the normal stress testing requirement.

For instance, what would cause your largest customer (or supplier) to fail? What would be the consequences for your business? What does this tell you about your business?

As with all risk management, it is up to the Board to use its experience and understanding about the way the business works to decide on the appropriate reverse stresses, what their impact could be and whether action is required.

Example 9: Aircraft safety

Aircraft engines are routinely tested to destruction, and aircraft are expected to be able to fly (and land) when one or more engines have failed.

In July 1989, United Flight 232, a three-engined DC-10, had an engine failure; the turbofan disintegrated and in doing so ruptured not only the main controls for the aircraft's tail and flaps, but the two backup systems as well.

Technically this rendered the aircraft unflyable, however the crew retained enough control using just the two remaining engines, to descend to an airfield and make a crash landing, saving 185 lives, although 111 were lost.

In this case the stress test might have been:

- What can cause an engine to fail?
- Can the aircraft fly on two engines?

The reverse stress test is:

- What functions does the aircraft need to fly?
- How might these be affected by an engine failure or explosion?
- How do you train the crew to cope?

6 Randomness and uncertainty

Before we consider some of the more quantitative methods of evaluation, it is worth briefly considering the nature of risk as it faces a corporation, especially from financial markets.

6.1 Statistics and markets

Much of the mathematics behind risk evaluation is based on concepts such as random walks, normal and lognormal distributions and so on. These, despite their seeming complexity, are actually quite simple to describe mathematically, and therefore make the ensuing analysis fairly straightforward. However, the risk manager must ask himself whether the real markets do actually follow these implied behaviours. Will any market follow a random walk from now on? As an example, consider a simple coin toss. In the short run, the outcomes are unpredictable but in anything beyond a few tosses, a pattern (for the split of total throws) emerges and is almost entirely predictable or certain in the long run, taking all tosses into account. Do financial markets really follow the same path? There is arguably no long run average for market rates and asset prices.

Ultimately markets are run by humans, subject to emotions such as panic, greed, herding, loss aversion and other biases and behaviours. Models are the best we can do in the circumstance.

6.2 Uncertainty

With a large portfolio, such as a large number of loan assets to diversified customer or a large number of insured customers, or a large number of loan losses, there is essentially no uncertainty. Risk management becomes not so much a matter of dealing with risk, but of ensuring that the portfolio is representative of the portfolio for which statistics are available. The approach is very much one of optimisation of portfolios.

Corporates do not have this business model, except perhaps in their manufacturing operations where every business has something they do in volume, cheaply. With financial risk, there is

often no portfolio. Every risk is essentially the first toss of the die when no long term average is there to stage a rescue. Evaluation of risk must take this into account.

7 Statistics

An understanding of basic statistics is essential for any treasurer because the language of the financial markets and particularly the investment markets has been worked around them. The major concepts that need to be understood include:

- random variable
- expected value
- mean (and other measures such as mode, median, geometric mean)
- variance
- standard deviation
- frequency distribution
- normal distribution (and simple calculations)
- lognormal distribution
- regression analysis and correlation

These concepts are dealt with in detail in Appendix A.

8 Correlation and portfolios

Correlation (how events are connected) underpins many aspects of risk management. Let us think again about the concept that in a large portfolio there is no uncertainty because the failure rate follows a defined frequency and the instances of loss are essentially unpredictable and uncorrelated. If there was a tendency of the portfolio to show copy cat losses, then the principle is undermined. Thus, if the loss events are uncorrelated, the portfolio is less risky than if they are connected. This is called the diversification benefit.

This is why insurers and bankers need large portfolios and why they try hard to ‘balance’ their portfolios with appropriately (un)correlated constituents.

Example 10: Lending portfolios
In the credit crisis, in the UK, several banks which specialised in property lending, as well as some building societies, were the ones that suffered most badly. Their assets were concentrated in property, where all assets were generally correlated and which all collapsed in value together. The portfolios were therefore more risky than a more diversified portfolio where the loans were made to unconnected business sectors.

In fact, the modelling shows that a portfolio of uncorrelated assets can show lower risk (standard deviation of returns) for a similar return. The Holy Grail of investment is to find assets whose returns are generally uncorrelated with others.

Correlation between assets, and of course risks, therefore becomes an important measure when considering enterprise risk. Treasurers must remember, however, several things about correlation. The first is that the mathematics can always be performed on two assets or risks to come to an answer for correlation, but that does not indicate real robustness in the result. Second is that correlations can be different over different time periods. There is much choice available when measuring time in correlation. The period of change can be as low as a second to as high as a year. The period of measurement can be from as low as a day to as long as 20 years or so. Choice here can give lots of different numbers. Finally, correlations can change over time.

It is also the case that in times of severe stress, there is a tendency for correlations to bunch. Thus in the credit crisis, many normally uncorrelated assets all fell in value at the same time and to roughly the same degree, thus destroying the diversification benefit.

9 Mean variance model

A building block of risk management for many years has been the mean variance model. It is the mathematical simulation of portfolios to see how collections of different assets (and liabilities) perform together.

The INPUT into the model for each asset or liability comprises:

- the expected (mean) return of that asset
- the volatility (standard deviation or √variance) of that asset
- the expected correlation of returns of that asset against other assets in the portfolio

The OUTPUT of the model is

- the expected (mean) return of the portfolio
- the volatility (standard deviation or √variance) of the portfolio

The model is therefore quite powerful and allows a risk manager to do several things. Firstly a portfolio can be optimised for risk and return. Thus the model can find what combination of assets will achieve a given return for the minimum risk. This is the model which makes investors seek out assets with returns uncorrelated with other assets as that is what drives portfolio risk down. Secondly the model will give a distribution of returns for the portfolio which allows a risk manager to find the probability of certain returns. This of course allows the concept of Value at Risk, where confidence levels can easily be found for whole portfolios.

Example 11a: Input into a mean variance model

	(of annual total returns)		Cash	ILGs (>5s)	FIGs (all)	FIGs (>15)	UK £ Credit (all)	UK £ Credit (>10)	Convention al Property	HLV Property	Hedge Funds	Hedged O/S equities	Unhedged O/S equities	UK Equities	Emerging/Pri vate Equity	Cash flow matching fund	Immunisatio n fund	Global HighYield / EM Debt	Liabilities
	Mean	Volatility	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
sh	4.5%	0.6%	100%	3%	10%	3%	11%	-1%	-14%	-12%	8%	6%	1%	6%	6%	3%	3%	9%	0%
s (>5s)	4.3%	8.0%	3%	100%	75%	75%	67%	69%	38%	62%	1%	32%	27%	11%	7%	90%	80%	3%	90%
s (all)	4.3%	11.0%	10%	75%	100%	97%	90%	92%	25%	61%	30%	30%	30%	30%	30%	90%	80%	-2%	77%
s (>15)	4.3%	11.0%	3%	75%	97%	100%	86%	90%	29%	68%	30%	30%	30%	30%	30%	90%	80%	-6%	80%
£ Credit (all)	5.1%	5.1%	11%	67%	90%	86%	100%	91%	26%	56%	0%	7%	-1%	-10%	-9%	12%	78%	9%	67%
£ Credit (>10)	5.1%	8.7%	-1%	69%	92%	90%	91%	100%	32%	64%	2%	16%	7%	-3%	-6%	8%	81%	5%	69%
nventional Property	6.5%	14.0%	-14%	38%	25%	29%	26%	32%	100%	88%	13%	58%	56%	60%	45%	10%	27%	7%	38%
/ Property	5.8%	12.6%	-12%	62%	61%	68%	56%	64%	88%	100%	5%	53%	46%	43%	31%	10%	53%	-2%	62%
dge Funds	7.5%	17.0%	8%	1%	30%	30%	0%	2%	13%	5%	100%	50%	50%	50%	50%	40%	50%	43%	10%
dged O/S equities	8.2%	16.8%	6%	32%	30%	30%	7%	16%	58%	53%	50%	100%	92%	82%	67%	40%	50%	19%	30%
edged O/S equities	7.6%	17.0%	1%	27%	30%	30%	-1%	7%	56%	46%	50%	92%	100%	81%	77%	40%	50%	18%	30%
Equities	7.6%	17.0%	6%	11%	30%	30%	-10%	-3%	60%	43%	50%	82%	81%	100%	74%	40%	50%	18%	30%
erging/Private Equity	11.0%	33.4%	6%	7%	30%	30%	-9%	-6%	45%	31%	50%	67%	77%	74%	100%	40%	50%	25%	7%
sh flow matching fund	4.6%	9.5%	3%	90%	90%	90%	12%	8%	10%	10%	40%	40%	40%	40%	40%	100%	90%	11%	100%
nunisation fund	5.1%	10.0%	3%	80%	80%	80%	78%	81%	27%	53%	50%	50%	50%	50%	50%	90%	100%	27%	90%
bal HighYield / EM Debt	6.0%	7.2%	9%	3%	-2%	-6%	9%	5%	7%	-2%	43%	19%	18%	18%	25%	11%	27%	100%	0%
bilities	4.3%	7.9%	0%	90%	77%	80%	67%	69%	38%	62%	10%	30%	30%	30%	7%	100%	90%	0%	100%

Example 11b: Output from a mean variance model

eficit at Risk = **$83** at the 95% confidence level **Wost case increase in deficit over 1 year =** **$65** at the 95% confidence level

ortfolio		1
ssets and Liabilities	**Duration**	**Value / £**
ash	0	64
ore Bonds	4.5	587
ong Bonds	12	0
igh Yield Bonds	0	64
arge cap equities		128
id cap equities		32
mall cap equities		32
nternational equities		64
edge funds con		0
ong Bonds (Mellon)		102
AA (Bridgewater All weather)		102
AA (Pimco All Asset)		102
pare		0
pare		0
pare		0
pare		0
iabilities	11.8	-1,475
wap hedge		
wap amount		550
ercentage liabilities hedged by swap		22%
ercentage liability duration hedged asset duration		45%
otal percentage liabilities hedged		67%

Risk decomposition

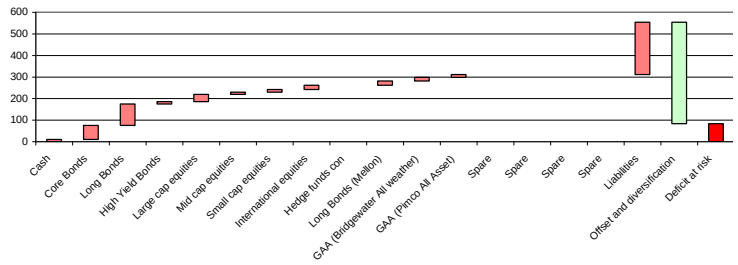

Risk Funnel

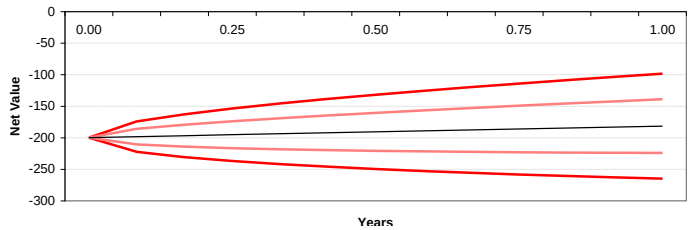

Example 11 shows the input to and output from a mean variance model looking at a defined benefit pension scheme. Notice the large number of inputs, reflecting the detail required for each asset. The output can be shown in many different ways, from simple expected return and variance, to Value at Risk, to a distribution of outcomes.

The mean variance approach models Modern Portfolio Theory, is very straightforward and easy to understand and the modelling is also relatively straightforward and indeed can easily be done on a spreadsheet. Being able to model things personally is a great benefit because it frees a risk manager from the clutches of bankers or consultants, speeding up the whole modelling process, and leaving the risk manager with a degree of control. If an outcome of an event or asset can be defined in the input terms, then this can produce powerful output.

The disadvantages centre on the ability to describe the behaviour of the component assets in the correct terms and perhaps especially in finding the correlation of returns of the component

assets. There is also a weakness in that the distributions assumed do not generally correspond to the distributions that we actually experience. In particular, real-life distributions tend to contain more extreme events than implied by the mean variance model.

The concept of Value at Risk has, as seen earlier, been heavily discredited in the financial crisis and partly because of the weaknesses explained here. However, perhaps we have pushed our use of the model too far beyond what it was capable of, and also relied on it too much – to the exclusion of other approaches.

10 Monte Carlo and stochastic methods

The advent of cheaper and cheaper computing power, combined with developments in various financial theories, have moved the mathematical approach to risk from the simple mean variance approach to ones that claim to be nearer to what really happens in markets and real life. The outputs that treasurers see in bankers' and consultants' pitch books when risk decisions are aired are almost certainly the result of Monte Carlo methods.

Monte Carlo methods rose to fame with the development of atomic weapons in and after the Second World War. Broadly they replicate what the mean variance approach does but with much more sophistication and versatility. They are based on the repetition of an event many times in an attempt to achieve a distribution of outcomes. Modern computers allow this very easily and some software programmes will allow the simulation to be embedded into spreadsheets so that the output distribution can be modelled and displayed. The inclusion of the distribution into models allows risk managers to investigate sensitive variables and concepts such as tipping points, i.e. sensitive points that can cause damage to a firm.

The INPUT into the model for each variable comprises:

- An assumed distribution for the behaviour of the variable
- Assumptions around correlated behaviour between variables
- The business model or portfolio construction

The OUTPUT of the model is:

- A distribution of outcomes

The advantages of a Monte Carlo approach over a mean variance approach are broadly:

- The flexibility to include different distributions, especially customised ones
- The ability to include the outcomes into models

The disadvantages of such an approach are broadly:

- Assumptions required may still be not a reflection of the real world
- The complexity in the model can be overpowering

An example of a simulation is useful to understand the power of the technique.

Example 12: Simple Monte Carlo simulation

Whistle Ltd has a business model whereby its one product is sold at USD 800 per piece and all the sales proceeds are converted into GBP at the prevailing spot rate. Its variable costs are EUR 500 per piece. Fixed costs in a period are GBP 4,000 and EUR 6,000.

Volume varies uniformly between 50 and 100 pieces, The GBP/USD exchange rate varies uniformly between 1.20 and 1.40. The GBP/EUR exchange rate varies uniformly between 1.1 and 1.3. Each of these three variables is assumed to be independent of each other.

The expected profit from this business model is as follows:

Sales	75 X USD 800 ÷ 1.30	=	£46,153.85
Variable costs	75 X EUR 500 ÷ 1.2	=	£31,250.00
Fixed costs	£4,000 + EUR 6,000 ÷ 1.2	=	£9,000.00
Profit		=	£5,903.85

This is modelled in the accompanying spreadsheet *Whistle Ltd Monte Carlo*. In the spreadsheet this is modelled approximately 1,000 times where each of the variables takes a value at random between the ranges indicated. Note that each time the sheet is opened, the random numbers are re-calculated and a different average value for profit is made. The sheet also calculates the standard deviation of profit and also the number of times that a loss is made.

It should indicate that the average profit turns out as expected, that the standard deviation is of the order of £4,000 and that losses occur around 4-5% of the time.

Use the basic approach shown to either adapt the model with different assumptions or make up one of your own. In particular think about how the two currencies might move together and how that might be modelled.

Example 12 is a simple example in that the distributions are uniform and the variables independent, but models can be designed to deal with the real complexities. It is important for treasurers to understand the assumptions and weaknesses that go into these models. It is rarely the maths that is wrong, but commonly the assumptions; and while 1,000 outcomes are modelled, only one will actually happen.

10.1 Stochastic modelling

It is worth commenting briefly on the term ‘stochastic’ because it often appears on the language of risk, and often from bankers and consultants. Stochastic simply means a description applied to systems whose behaviour is intrinsically non-deterministic, i.e. in some way random. Thus, a system's subsequent state is determined both by the process's predictable actions and by a random element. Also any kind of time development (be it deterministic or essentially probabilistic) which is analysable in terms of probability deserves the name of stochastic process.

11 Presenting the data

We have seen several different methods of approaching the problem of evaluating risk and suggested several different measures that could be adopted. The problem next facing the risk manager is to present the output of these techniques to management or risk committees and to the board. This is difficult and presentation should be borne in mind when considering the evaluation techniques.

Exhibits 6 and 7 are examples from presentations derived from different risk measurement techniques.

Exhibit 6: Risk measurement output

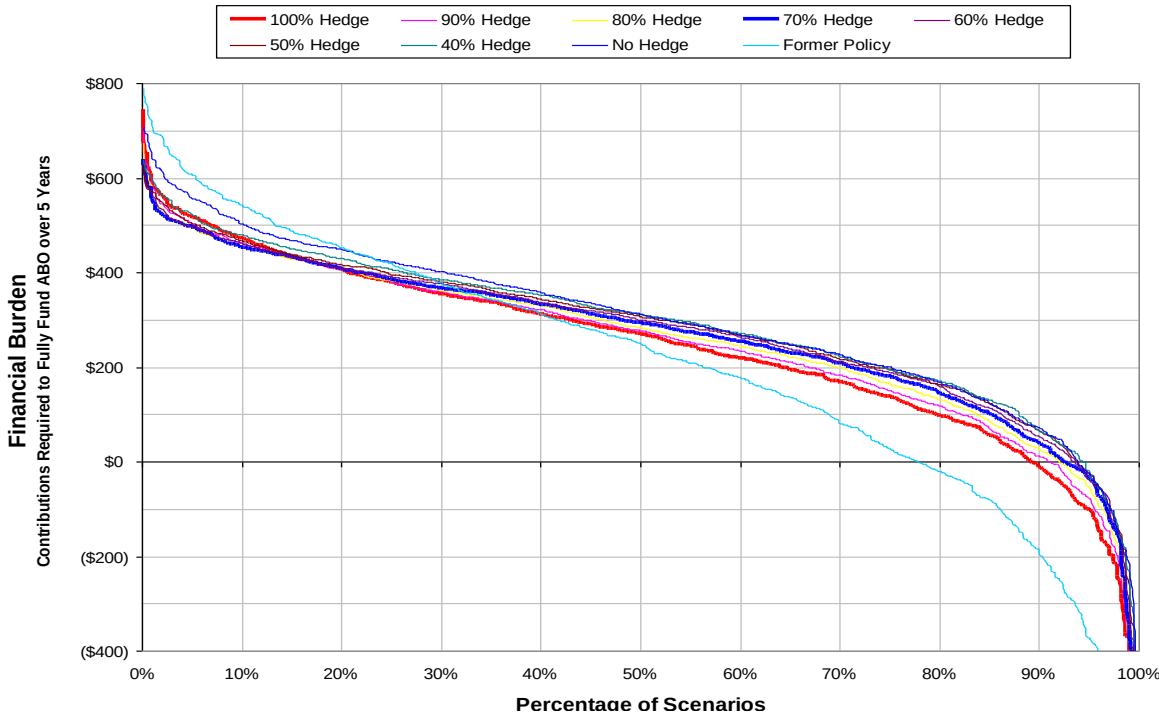

Exhibit 7: Risk measurement output

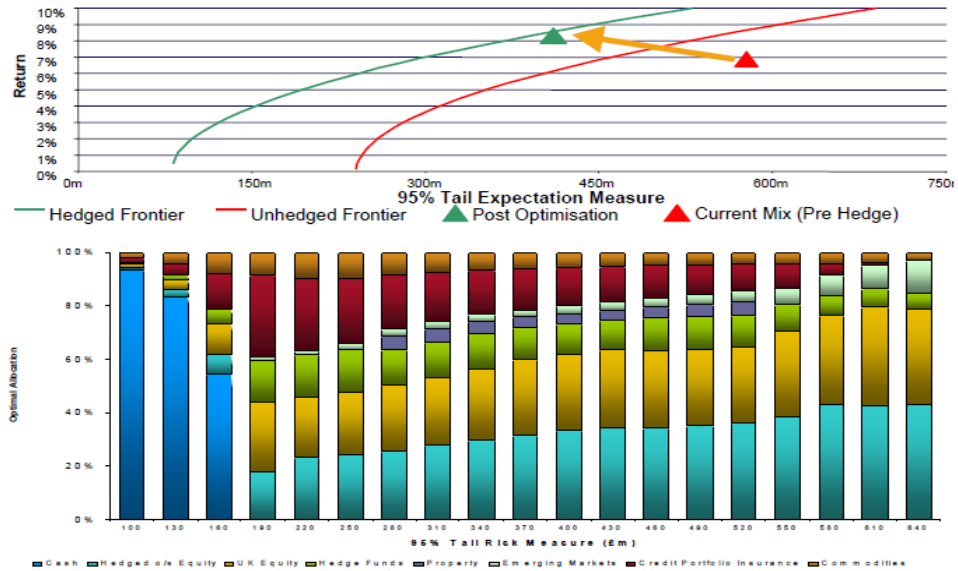

Without considerable immersion in the particular subject and the measurement techniques and assumptions, it would be a challenge for most people to understand quickly what is going on in either of these graphics or indeed in many presentations shown to committees or management. This is no criticism of the models involved or indeed the technical outputs and to those specialising in risk management the graphics are extremely useful.

However, there is a risk of such presentations being ignored if the recipients cannot quickly understand them. Such presentations are likely to form only part of a much larger set of documents and so can quickly be glossed over.

A second issue is to try and work out which scenarios to present to management. For example, a proposal presented to management to fix a certain proportion of interest rate bearing liabilities might include the ratios of say 40%, 50% and 60%, which is the range recommended by the treasurer and the presentation will include all the advantages and disadvantages in some way. Management may rightfully query this and say “What about 80% or 100% fixed?” Working out the answer to this might be quite difficult when put on the spot and the chance to persuade management will have passed by. The treasurer must be prepared! It is probably unlikely that such questions will be asked as the treasurer will be trusted to propose only sensible ideas. However, management is perfectly entitled to, and arguably should, pose such questions.

This is an example of the power that lies in the hands of those presenting data. They can always present things in a way to advance their own position by selecting the data very carefully. This is a particular issue when banks make presentations in support of their own products.

12 What can go wrong?

The nature of risk is that events will happen which have adverse consequences for a firm. Returns above the risk free rate are only achieved at the risk of such consequences. This portfolio approach leaves investors with a diversified portfolio relatively well protected so that losses are offset by gains, but it also means that some individual investments, firms where risk managers are employed, will fail or perform poorly although of course some do better. Failure also does not mean that risk evaluation failed. If risk evaluation identified a major risk, then there is a possibility it will happen.

The major causes of failure of evaluation include:

- poor quality input data (rubbish in / rubbish out)
- incorrect correlation assumptions or changes
- risk management performed in silos
- different measures of risk involved
- tail risk

12.1 Correlation measurement errors and changes

Most measures of risk where portfolios are involved, (i.e. any more than one variable), make assumptions about how the components of the portfolio, or the variables, behave together. This behaviour underpins portfolio theory and diversification, but the theory depends on the correlation of the assets remaining identical over the life of the risk measurement. Should this fail, then the whole risk evaluation is potentially in error.

Example 13: Correlation change

Hunter AG is a diversified multinational with two divisions, one of which operates in the basic food market, the second operating in luxury goods. Traditionally basic foods have performed well in a downturn, as people resort to basic foods, but poorly in good times, when more exotic foods are consumed. Luxury goods is the opposite, performing well only in the strong economic times. Management consider their correlation to close to -1 (exact opposite) and indeed overall profits have proved very static.

Following the financial crisis of 2007, the basic food business has performed well as expected but the luxury goods market has also performed well. The rich seemed to have remained rich, or unwilling to make sacrifices, and the luxury market consumers seem to have kept up expenditure. In addition, support from the growing markets of Russia and China have kept demand ahead of expectations.

The correlation between divisions has changed from -1 to near +1, on this occasion to the benefit of Hunter AG. However, it seems that the Eastern markets are at risk of a downturn at the same time as the basic food business might also turn down.

Hunter has gone from being low risk to very risky.

It is possible that we have pushed the modelling too far. It is simple mathematics to extend the time period for most risk models, but time is a most powerful agent of change, undermining such models' reliability.

12.2 Risk management in silos (enterprise risk management failure)

The nature of management is to delegate and this applies to risk management as much as any other type of management. It is inherently attractive to carve up risks to the relevant manager and indeed treasurers would be rightly shocked if they were not heavily involved in the financial risks of the firm.

However, risks must be considered as a whole. If, for example, interest rate risk is not considered within the context of the firm, then there is a danger that risks are multiplied.

Example 14: Interest rate risk

High Rise Inc is a property company and manages risks individually. Rents are managed by the sales and letting department in accordance with market pricing. Debt and interest rate risk is managed by the treasury department. In accordance with suggestions from banks, High Rise has adopted a policy for interest rates whereby fixing is kept to a minimum (short term rates almost always being cheaper in the long run than long term rates) but some 'out of the money' cap protection is taken.

Interest rates rise after a long period of low rates and the interest charge rises above rental income, which has not increased significantly.

High Rise thus risks failure for a fundamental failure of its business model where the interest rate risk was not compared to the underlying business risk. Clearly High Rise had high financial risk (high leverage) which should have been reduced.

12.3 Different measures of risk adopted

Even if Enterprise Risk Management is followed, then unless a common measurement is adopted, there is no way to consolidate risks and work out what is most important to the firm.

Example 15: Different measures of risk adopted

Follower SA operates Enterprise Risk Management and risks are measured and brought back to a central risk committee. The operating management report in terms of sales impact, i.e. percentage of sales at risk in different geographies. There is no indication of margin or working capital implications of lower sales volumes. The treasurer is most concerned with performance against bank covenants which are maintenance covenants measured against interest cover and Net debt / EBITDA.

There is no connection between the two risks at the risk meetings and it is not possible to say how performance against covenants is impacted with a mere percentage of sales forecast.

The treasurer accordingly requests that the percentage drop in sales at risk be fully modelled into forecasts around the bank covenants.

12.4 Tail risk

The distributions often adopted in risk management are typically normal or lognormal distributions – they involve simple mathematics which can easily be manipulated to achieve sensible measures. These distributions are what options pricing is based on (Black Scholes). However, experience tends to show that the extreme ends of the distributions of actual events do not follow this pattern.

Exhibit 8: Fat Tails

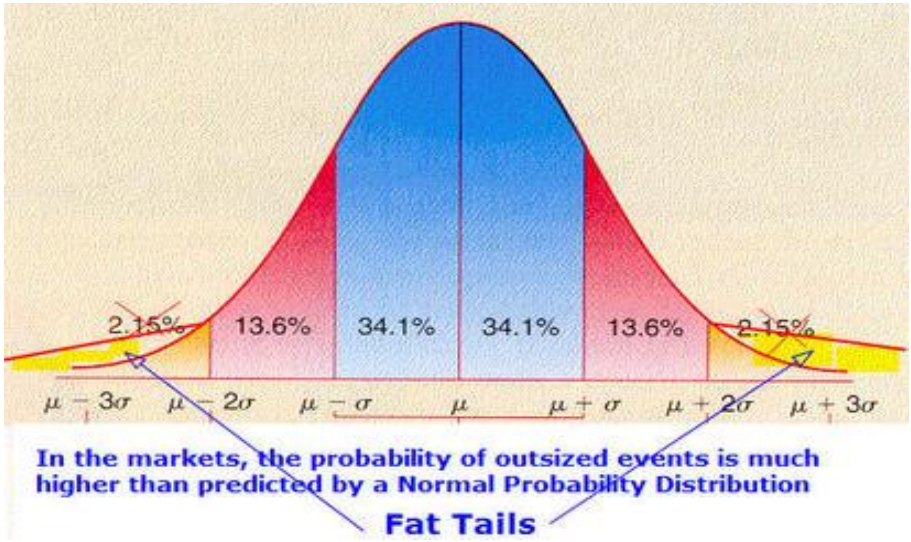

Exhibit 8 shows a normal distribution with the real experienced Fat Tails superimposed. Thus extreme movements are more common than is theoretically expected. Modern risk management is starting to tackle this although it is interesting that options pricing in practice does incorporate this increased tail risk, the so called 'Volatility Smile'.

Exhibit 9: Volatility Smile

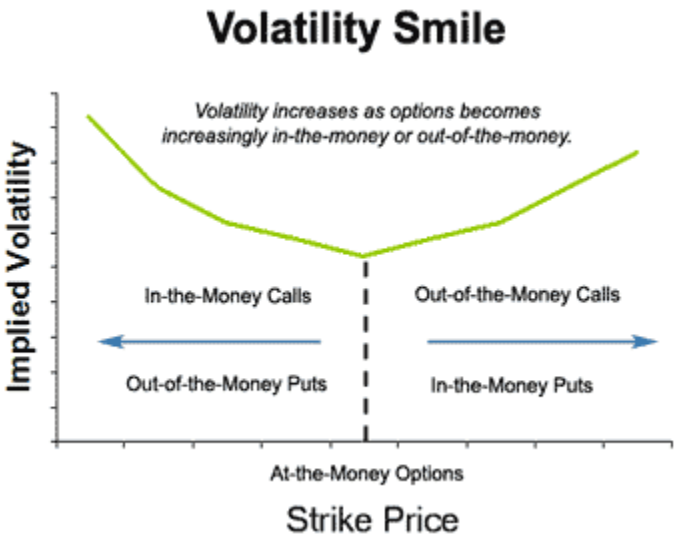

The options market, trading as it does on volatility, is of course a market in the price for risk and is a very useful way to approach the measure of risk. The implied volatility allows calculations to see what moves the market might be expecting.

13 Trusting the market

Financial risk is generally market based. To a great extent treasurers are lucky in that many of the risks encountered have liquid and transparent markets where that risk can be hedged away in some fashion. It is worth thinking briefly about the treasurer's interaction with those markets. Several aspects are worthy of mention:

- Markets carry tremendous amounts of information.
    - Movements in prices or implied volatility happen for a reason.
- Markets can be very wrong and overshoot in any direction.
- Markets suffer from occasional poor liquidity and so have periods of high volatility.
- Markets are not a good indicator of future prices.
- However extreme prices might be, they can always become more extreme.

Despite these factors, it is still the case that the markets are the best repositories of information and to argue that the market is wrong is fraught with danger.

Exhibit 10: Sterling interest rate forwards and actual outturns at mid 2011

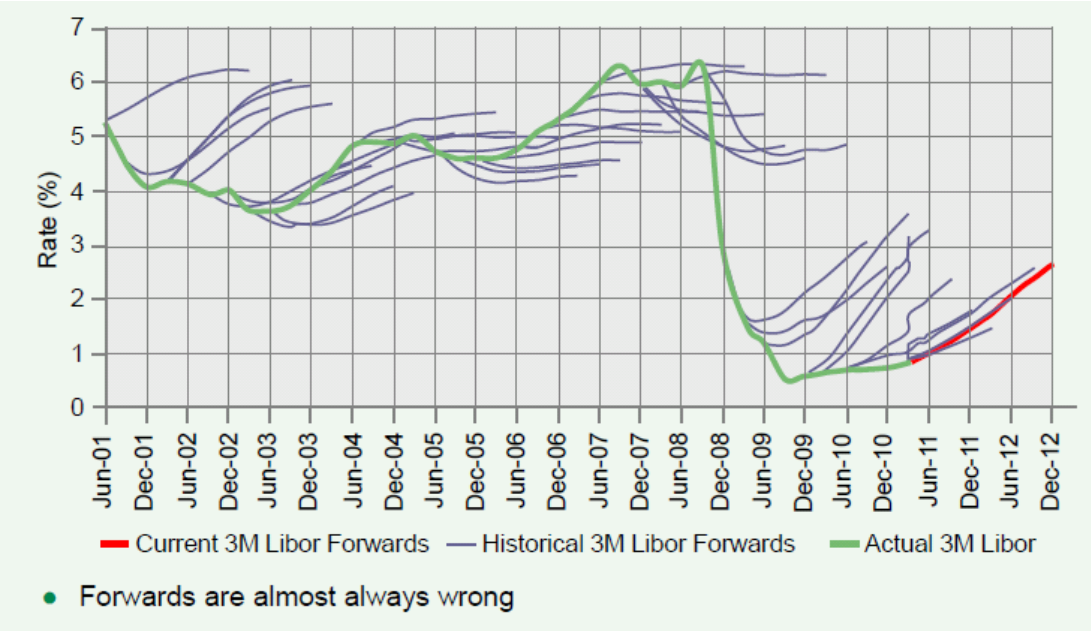

Source: Bloomberg and Lloyds Bank

Exhibit 10 is an example of one market where forward pricing has been consistently 'wrong' (in that the future price available today has been substantially different to the price available in the future). It might be argued in this case that eventually the market does go where it is expected to go, but with no degree of reliability. However, all the information available at the time of the future pricing is incorporated in the market *at that time* as anticipated by the Efficient Markets Hypothesis (EMH).

However:

- The EMH incorporates all information, including expectations, not necessarily the 'right' information; and
- As time moves on, new information comes to light and previous expectations turn out in fact – both are incorporated in the price, so no wonder the actual future price is different.

A treasurer cannot be expected to hit the peaks and lows in markets but perhaps should have some insight into when markets are oversold or overbought. This is difficult territory though, because any opinion on the state of the market should always invite the question "Why is the market wrong now?" The treasurer must have an answer to this question if a decision is based on such a view. It takes a lot of effort to follow a market and reach conclusions on which to make decisions and treasurers may find it difficult to devote sufficient time for any one decision.

Example 16: Mr Market
Benjamin Graham's 'Mr Market' analogy is apt[2]. Imagine someone who, every day, knocks on your door and offers you a price for something you own. Sometimes you will perceive the price to be ridiculously high, and sometimes it will be ridiculously low. It is up to you to decide the value of your asset, and sell it when Mr Market offers you a higher price rather than a lower one. Perhaps you will decide that the asset is cluttering up the house and anyone who will pay you for it is welcome to it; perhaps you will decide that its sentimental value means that you will never sell at any price. But no matter how many times you reject the offer, Mr Market will always be back tomorrow with another price.

14 What you can hope for in evaluation

Arguably evaluation is so difficult that it is probably not worth the effort. Many treasurers take a similar view of cash forecasting, especially short term cash forecasting. However, without an appreciation of the scale of the risk, it is impossible to devote the right resources, from actual (equity) capital to time and effort in managing risk. The most that one can hope for in evaluation is to get a reasonable scale on the risk. In the same way that salesmen need to understand the margins on their individual products (so as to know which ones to push), so risk managers need to understand the scale of different risks to know how much capital to allocate.

The techniques indicated here are all useful at giving a scale, or order of magnitude, to risk. Practically, strict evaluation might be more difficult. A firm with several different business models and as many foreign exchange flow patterns will find evaluation very challenging.

Some other ideas include:

- Use more than one technique in order to provide a wider perspective
- Use judgement when interpreting numerical outputs. In all cases, does the statistical result make commercial sense?
- Consider a relative approach rather than absolute. So long as consistent measures are used, something that doubles a risk surely makes it twice as important.

We should also consider how risk evaluation interacts with investment appraisal methodologies such as discounted cash flow (DCF) analysis to assess whether a proposed investment will meet the return requirements of the firm. A DCF analysis will typically show a set of cash flows expected from a project which are then discounted at the firm's cost of capital to see if the net present value is positive. However, this process may take little account of risk if:

- the cash flows are just one outcome from a possible range
- the cost of capital used fails properly to reflect the risk of the project

[2] The Intelligent Investor by Benjamin Graham 4th Revised Edition 1973 but still in print from HarperCollins

Risk management tools can help here by assigning a range of outcomes to a particular project, perhaps with statistical analysis or even just sensitivity analysis. Risk managers can assess the riskiness of the cash flows, perhaps with scenario analysis or statistically-based methods if financial market variables are relevant. Different return requirements (i.e. costs of capital) may be applied according to the riskiness of the project.

This same sort of approach can be applied to normal budgeting and forecasting procedures, where often only one possible outcome is ever modelled by firms. Ranges of outcomes are a more realistic approach to the real world.

A proper response to risk, the next stage in the framework, is not possible without evaluation.