The Inverse Risk Logic Approach to Strategic Risk Management

By Dr David Bobker

Introduction

This article summarizes the approach taken in the Risk Management Module of the FIDE programme (Module 2) in which the role of the top down approach, called "inverse risk logic" is emphasized in strategic risk management. While Malaysian banks in general have fared well throughout the financial crisis which broke in August 2007, now is not a time for complacency. Indeed it may actually be the time for even more vigilance. History has shown time and again that crises are invariably preceded by "good times" during which attention to the growing risks is lost until it is too late to do anything about it. In fact our thesis is that the time for the greatest focus on risk is precisely when things appear to be going well so that the board can ensure that "profits" are real and will not later prove to have been the illusory profits of the crisis which quickly evaporated when the markets collapsed.

People have regularly asked about the role of "risk management" during the crisis. Many failed institutions apparently had near state of the art risk management in terms of computational systems, people and frameworks. For such catastrophic failures to have occurred, either the risk management frameworks and methodologies were in some way deficient; or they were not properly applied (for example due to governance failure); or there was some combination of both of these. It is our view that both governance failure and methodological failure were to blame. Much has been written on the subject of the former which relates to boards of directors being unable to exercise sufficiently strong oversight over executive management. However, it seems that there is still little challenge to the "value at risk/economic capital" paradigm on which most banking risk management, and indeed the whole Basel regulatory superstructure was built. This is in spite of the fact that the major shortcomings of this approach have been well known for many years. For example, in the introduction to a technical paper prepared in 2006[1] it was stated: : "…there is a concern that the vast global derivatives market, the number of unregulated hedge funds, the merging of financial markets across national borders and the explosive growth of private equity funds, make the financial system more unstable and susceptible to meltdown. These concerns are not new but have been serious topics of discussion for several years. The extraordinary fact however is that the volatility of financial markets today is about as low as it has ever been. This has been true for most of the years 2004- 2006." This precisely states the issue: namely, the time when growth in the markets was high and volatility as low as ever was just the moment before calamity struck. Economic capital models built on a foundation of market and default data collected during abnormally benign times were inevitably bound to give a hugely optimistic view of the risks building within the system.

[1] Colacito and Engle 2006 Term structure of risk, the role of Known, and Unknown Risks and Non-stationary Distributions

The objective of this article is to present in non-technical terms an overview of the accepted economic capital methodology, explain the limitations and suggest an approach to supplement (not replace) it in the risk and capital management processes in the board.

Note: the methodology is applicable to any organisation not just banks or financial institutions. We have dwelt rather more on banking VAR and economic capital as that was a key issue for the banking collapses. However, the same philosophy of economic capital underpins insurance regulation. But the idea of the critical vulnerability review is of universal applicability and rigorously applied might have been of benefit to a non-financial sector company like BP for example prior to its mega-disaster.

The Value at Risk Paradigm

Put simply, value at risk is an estimate of the worst loss that could occur, within a specified timeframe, to a specified probability. For example, the popular "daily value at risk" measure, often quoted in many banks' annual reports, is the level of loss the bank believes would only exceeded in one day with a probability of 1%. The 1% is a policy driven value. In fact any value could be used. For example attention was drawn, during the FIDE Risk Management Module 2, to a disclosure in Barclays plc annual report 2008 where they decided to change the DVAR probability from 2% to 5%.[2]

Defining "confidence level" as 100% minus the probability, DVAR could also be expressed as "having 99% confidence that market risk losses on any day would not exceed the DVAR value". The VAR paradigm is then based on the idea that by setting a DVAR limit the risk is managed to a tolerable level.

There is, however, a number of (well known) major flaws in this approach:

- Ensuring the DVAR of a portfolio is acceptable gives no useful indication of the probability of far greater losses. In other words, one may be 99% confident that DVAR will not be exceeded, BUT, there might actually be a 0.5% probability of a catastrophic, terminal loss;

- By choosing a time period of one day (the "daily" element of DVAR), there is an implicit assumption that loss making positions could be closed out within a day so that the level of loss could at least be limited. Unfortunately as was seen in the crisis, during stressed market conditions, market liquidity might "evaporate".[3] In other words, forced selling of illiquid contracts could result in far greater losses than the DVAR and may not even be achievable within

[2] Barclays plc Annual Report 2008 page 120.
http://www.barclaysannualreport.com/ar2008/files/pdf/Annual_Report_2008.pdf

[3] The full blown crisis is usually deemed to have begun on 9 August 2007 with a statement by the French bank BNP Paribas which froze two sub-prime funds to investor withdrawals. It said: "The complete evaporation of liquidity in certain market segments of the U.S. securitization market has made it impossible to value certain assets fairly, regardless of their quality or credit rating". BBC economic editor Robert Preston described the statement as "scary". The failure of Northern Rock occurred about a month later.

a day, or, as was the case of the "toxic" sub-prime CDO's may not be possible at all. . In any case for banks actively trading, the majority of contracts might be OTC in which case closing them out actually involves finding a counterparty and doing an equal and opposite trade which again may not be possible, and would increase credit risk exposure even if it were;

- Another problem with OTC contracts is valuation. Especially for more complex "structured products" which are not exchange traded, valuation of the contract depends on a model (these contracts are "marked to model"). A model depends firstly on the assumptions behind it and secondly on data inputs. Such assumptions carry a serious risk of becoming invalid during market stress. For example, models of contracts with an element of "optionality" depend on an estimate of volatility since it cannot be observed directly in real time. The models themselves are also likely to depend on a number of standard finance theoretic assumptions which are known not to be empirically correct and have often been observed to have broken down completely during periods of market turbulence.

- The calculations are normally done based on data collected in the near past. This is the so called "driving while looking at the rear view mirror" problem. In order to do the risk calculations, statistical analysis of market data is used to build the various models. By restricting it to the recent past, there is an implicit assumption that future trends will follow the same probability distributions as in the recent past. In other words, the nature of the risk is static. Extending the time period is normally ruled out because old data is considered to have limited significance today. This is why the regulations insist on "stress tests" to supplement the statistically based calculations. We return to stress testing later.

Economic Capital

Economic capital then takes the idea, originally developed for market risk, further and applies it to the whole institution. Under Basel 2 regulation the idea was to try to ensure that a bank complying with the capital rules would only have 0.1% probability of annual losses exceeding its regulatory capital base. 0.1% is equivalent to a return period of once in a thousand years and was considered a more than reasonable level of probability of bank failure.

Once problem with the approach is a general misunderstanding of the meaning of 0.1%. More than one financial institution director has said words to the effect that "I'm not going to be around in 1000 years so there's not much to be worried about!"

But far greater difficulties emerge once a bit of thought has been devoted to working out what exactly does the 0.1% mean? We will examine that issue below, but for now we examine another issue relating to the 0.1%. As stated above, the 0.1% refers to the supposed probability of all the bank's capital being lost in a year. But we would contend (which was demonstrated admirably during the 2008 meltdown following the bankruptcy of Lehman Brothers) that a bank need only lose a proportion of its capital base to become unable to continue its operations and go out of business. We term that point as "non-viability" and it will happen to a bank, where confidence is everything, long before it has burned *all* of its capital

A good working assumption is that if a bank's risk asset ratio falls below the absolute minimum of 8%, prescribed in the rules, then it would become non-viable as an independently operating entity. To take a simple example, suppose the bank has a capital ratio of 15% (that is, regulatory capital is 15% of risk weighted assets). Then a loss of 7% of RWA will lead it to just about breach the minimum level. But that 7% translates into 7/15, or 47% of its regulatory capital. In other words, the bank will go out of business after a loss of 47% of its capital. The big question then is this: if the probability of losing *all* the regulatory capital is 0.1%, what is the probability of losing 47% of the regulatory capital? Clearly it will more than 0.1%, but the calculation of loss probabilities is highly non-linear which means that non-viability may be reached with a probability of 1%, or 5% - one cannot say without looking at the risk curve used to assess the economic capital.

But even if one could, we still have all of the points raised above in relation to market risk VAR. Relying completely on that 0.1% number to ensure a negligible probability of failure of the institution is clearly not a wise strategy given the myriad assumptions implicit in the calculation and all the uncertainties which surround them. For example, in a classic study[4], different risk software vendors' products were used to compute VAR associated with the same specified portfolio. The results were surprising in that the different "calculation engines" came up with startlingly different answers. Yet the board may be being asked to effectively "bet the bank" on the basis of such unreliable numbers. The old adage – beware spurious accuracy applies very forcefully here.

What is Probability?

Much of the quantitative consideration of risk depends on assessments of the likelihood of an event occurring. This would normally be expressed as a probability, i.e. a number between zero (impossible) and one (certain) which expresses the degree of uncertainty. However, probability in its very nature is subjective. In fact the definition used by so called "Baysian" statisticians, actually defines probability as an individual's "degree of belief" that an event will occur. Whilst in essence subjective, there are many cases where people will tend to agree. For example, people will naturally assign a probability of 0.5 to the event that the toss of a coin shows heads. But why is this so? Might it not be different if there were a lot of money at stake and you strongly suspected the person tossing the coin to be crooked? In fact our degree of belief may be stronger or weaker depending on how much prior knowledge we have about the mechanisms which are responsible for the final outcome. Thus assigning a probability to "it will rain tomorrow morning" is likely to provoke a much wider range of assessments - and of course none of them can be wrong unless impossibility or certainty were assigned. This is why assigning probabilities to events which may terminally damage an organisation can be so dangerous. The basic problem is that quantitative risk theory is based on an (implicit) assumption that one is working on a problem similar to the coin tossing experiment. The basic assumption is that there exists an (unknown), but static probability of the coin landing heads which makes the problem highly tractable and susceptible to classical analysis. Unfortunately the reality is that the unpredictability of the behaviour of people, and markets is such that the problem lies much more in the domain of whether or not it will rain tomorrow type of assessment

[4] Pritsker, M 1997 Evaluating value-at-risk methodologies: Accuracy versus computational time

and one man's word is as good as another's. Therefore the statistical models are bound to give a far greater sense of security of knowledge than is warranted.

The bottom line for directors and senior managers of banks and insurance companies is this: "calculations" of probabilities are merely working assumptions based on our current knowledge and assumptions about the future can never in any sense be "correct". Put another way, **VAR, and economic capital, cannot be a measure of risk they merely represent best estimate working assumptions. Basing a strategic risk assessment and capital adequacy on such statistical models is a big risk itself.**

A Simple Game

To illustrate, consider the simplest probabilistic experiment, namely, tossing a coin a number of times and recording the results as a series of "H" or "T" results according to whether the coin lands heads or tails face up. The outcome of the game is the number of times the coin comes up heads and this may be used as the basis of a bet.

Most people when asked will say without hesitation that the probability of a coin landing heads up is 50%. In workshops we ask the following question: suppose you toss the coin ten times and it lands heads up on nine occasions – what is your view of the probability now? The majority stick to their guns and say 50%. One or two (usually out of 10 – 15) start to look a little suspicious and say nothing. We now ask, suppose the coin comes up heads 99 times out of 100, what is your view of the probability now? Many people continue to hold out with their 50% view at which point they are asked "what would it take to convince you the answer can't be 50%?".

A simple calculation[5] reveals it highly unlikely that the probability of heads is 50% and the most likely value is 90%. Baysian statisticians would do it slightly differently and might start with a clean sheet, assuming the probability of heads could equally likely be anything between 0 and 1 and modifying their view in the light of the data. Either way there is a feedback loop which gradually modifies our view of the probability as more data is gained.

Unfortunately even that will not work for banking risks. The difficulty is that we are not dealing with a static but unknown probability which we gradually estimate with ever more accuracy as more data is obtained. The point is that, using the coin tossing analogy, the probability of heads is *not static* but moves in unpredictable ways. Thus even a continual re-assessment of its current value will not help us calculate the risk.

This has big implications for risk management

[5] **IF** the probability of heads is 0.5, **THEN** the probability of nine heads coming up in ten throws is 10/1024 or less than 1%. A statistician would reject the hypothesis of 50% and would carry out a *maximum likelihood estimation* to conclude the most likely value of the probability of heads were 90%, not 50%. Similarly on the basis of the hypothesis that Prob(Heads)=50%, then the probability of of 99 heads in 100 tosses is $100/2^{100}$ which is effectively zero. The most likely value for Prob(Heads) is 99%

The principle of retail lending (and the principle is exactly the same for retail insurance) is that by the laws of large numbers, the deviation from the "expected value" becomes progressively smaller as the number of items in the portfolio increases. Suppose we toss the coin 1,000,000 times. The "expected value", for a fair coin is then 500,000 heads. However the risk manager can easily calculate the whole risk curve and he or she would say that the VAR to 99.9% confidence is 501,500[6] and the probability of 600,000 was as close to zero as makes no difference – an impossibility in other words.

Let us now suppose that the equivalent of world financial market turmoil ensues and the outcome *is* in fact 600,000. What do we think now? Well we reject the hypothesis that coin is fair, and using our feedback loop we reassess the probability of heads for each toss as 60%. The only problem is we are now out of business facing myriad lawsuits, being asked to testify before the equivalent of Congressional Committees, and worrying we may face a jail sentence of many centuries (the US does not like to under-sentence).

The lesson from all of this is very straightforward: **the risk to the organisation is not that the VAR on the risk curve will be exceeded but that in fact the risk curve itself is wrong. It has morphed into a different curve and pulled the rug from under our careful planning based on meticulous statistical analysis before we had a chance to react.**

The truth of the matter is that the major weak link in conventional risk management thinking is the reliance on probabilities calculated from statistical analysis of data. However when circumstances change, so do those probabilities. During the crisis, a well known market risk expert said that the problem was that one in one million year events had been recurring about every six months.

So What Can Be Done?

This final section now examines the options available to boards in the light of the above arguments. This will relate very much to a consideration of **risk appetite.**

We reject the economic capital model which effectively says that if we can "calculate" our VAR to 99.9% confidence, then if we have a capital base of at least that, we only have a one in a thousand year probability of insolvency. To summarise our objections:

1. There is a much higher probability of *non-viability* than 0.1% which may give an unreasonably optimistic view;
2. Given the myriad assumptions and hypotheses, models and calculations which are required to calculate the economic capital it is very unwise fully to base important strategic risk decisions on such an intrinsically unrealiable figure;
3. The figure is based on a combination of statistical analysis of past data and "expert opinion". The expert opinion is clearly highly subjective and dependent on the particular expert(s) giving it. Secondly, basing a risk assessment on historic information implicitly assumes that

[6] The distribution is (approximately) normal with mean 500,000 and standard deviation 500. The 0.1% point on the normal curve is about three standard deviations from the mean.

the past is a good predictor of the future. Indeed the past *may* be a good predictor of the future, but only up until the time that it is not – which is precisely the risk!

The world's banking regulators implicitly agreed with this stance by requiring banks to carry out ***stress test.*** After the East Asian Crisis of 1997 Bank Negara introduced a rigorous stress testing regime on all banks as a key element of supervisory policy. Basel did not follow suit until much later with the introduction of Pillar 2 of Basel 2 to supplement the calculations of Pillar 1. But the reality is that for the last couple of years, regulators worldwide have been relying *almost exclusively* on the stress tests as a measure of capital adequacy and do not place undue reliance on the VAR and economic capital models. Indeed the current changes (known as Basel 3) propose introducing a leverage ratio limit which would be an additional, high level, top down limit. The new, greater reliance on stress testing is eminently sensible and is, with a slight twist, what is advocated here. Of course it begs the question: if the regulators base their real assessment of capital adequacy on the stress tests, what exactly is the point of all the rest of the 347 pages of Basel 2 calculations? That is a very good question which we hope to answer in a subsequent article. For now suffice it to say that all the billions spent on IT systems and expensive risk management departments have not been entirely wasted, but so far, too little of their output has properly directed towards *strategic* risk decision making. In other words, credit scoring systems, portfolio simulations, VAR "engines" and the rest are essential for competitive advantage and effective operational efficiency and effectiveness but cannot, on their own, be used for the capital calculation. In fact, the reverse stress testing approach we advocate requires use of those models but in a different way than is usual.

To return to the initial question "what can be done", our suggestion is to use the inverse risk logic methodology as the primary guide to strategic risk assessment and to ensure the board understands its own risk appetite and is in a position to take early action at the first hint of the organisation moving beyond it.

Inverse Risk Logic, Risk Appetite and the Critical Vulnerability Review

Inverse risk logic simply refers to the process of determining all possible scenarios that could lead to failure of the organisation and making the strategic decision as to whether the status quo is within the risk appetite of the business or whether a change of strategic positioning is required. The process of determining the potentially catastrophic failure scenarios is termed ***the critical vulnerability review.***

Although "risk appetite" is mentioned in every corporate governance code, Basel 2 and risk management guide, it is something for which as sensible definition is notoriously difficult to find. For example one risk "standard" "defines risk appetite more or less as the amount of risk the organisation is prepared to tolerate – hardly helpful.

In the past we have used a more quantitative definition, namely, "risk appetite is the acceptable probability of a catastrophe" which assumes that the board of directors has carefully decided what level of impact of an event would be "catastrophic" for the organisation. Such an exercise is in any

case highly desirable for the board to carry out as a necessary pre-requisite to formulating its risk management policies.

However, in the light of our strong scepticism of the value of the estimates of probability coming out of the normal models, coupled with the fact that much of the necessary input is highly subjective anyway, we would now put forward the following as a working definition:

Risk appetite is the willingness, or not, to tolerate the currently identified set of catastrophe scenarios each of which, if it occurred, would lead to the inability of the organisation to continue as an independent going concern.

In other words, we have fused the concept of risk appetite into that of assessment of probability since both are, at least in the context of strategic risk for financial institutions, subjective.

We may also add that this seems to be the approach of the legendary Warren Buffet and long time business partner, Chalie Munger. The latter, for example has been quoted as saying "all I want to know is where I'm going to die so I'll never go there". More simple and to the point than the above definition perhaps. However, the method we put forward here is perhaps explained more clearly in one of Buffet's famed annual letters to shareholders, the one published in Feb 1997 relating to 1996.

The following extracts from the letter explain the approach:

"We took on some major super-cat exposures during 1996. At mid-year we wrote a contract with Allstate that covers Florida hurricanes, and though there are no definitive records that would allow us to prove this point, we believe that to have then been the largest single catastrophe risk ever assumed by one company for its own account..... So what are the true odds of our having to make a payout during the policy's term? *[referring to re-insuring the California Earthquake Authority]* We don't know - nor do we think computer models will help us, since we believe the precision they project is a chimera. In fact, such models can lull decision-makers into a false sense of security and thereby increase their chances of making a really huge mistake. We've already seen such debacles in both insurance and investments. Witness "portfolio insurance," whose destructive effects in the 1987 market crash led one wag to observe that it was the computers that should have been jumping out of windows."

Buffet was taking on a gargantuan potential payout, actually a whole portfolio of them, so what was the thinking behind it? He goes on to say: "Large as these coverages are, Berkshire's after-tax 'worst-case' loss from a true mega-catastrophe is probably no more than $600 million, which is less than 3% of our book value and 1.5% of our market value. To gain some perspective on this exposure, look at the table on page 2 and note the much greater volatility that security markets have delivered us."

In other words, rather than putting faith in a calculation of the probability of occurrence of the worst case scenario, he simply decided that it was a bet he could afford to lose, but which would produce a great profit if he won and it is well known that General Re insurance does not come cheap!

That then, in a nutshell is the essence of the method. In the first place **the board must work out just what it is betting on and decide (a) if it can afford to lose and (b) whether the potential profit is worth the risk.**

Of course that simple advice is easier said than done and in fact part (a), deciding what the current bets are, is necessarily complex and specifically related to the organisation's business activities. If however, risk management departments start to orientate themselves more to providing clear information to the board so that they can make the appetite decision based on harder facts, the better will risk management be. **Above all, if the board cannot really understand what exactly it is betting on, better to stop the activity altogether than hope it will all be alright.**

Put another way, the inverse risk logic approach means calculating how far away from the abyss is the institution and deciding, judgementally, whether or not it is far enough. If the institution is dealing in products such that you cannot actually say how far from the abyss you are, best move back from the edge!

Why do we term this "inverse" risk logic? Simply because the usual approach to stress testing and "enterprise risk management" starts with the questions: what can go wrong and how likely is it? That leads to a list of potential risks which is so big that the big risks almost inevitably get drowned in a sea of paper containing highly subjective, generally superficial assessments and the focus on the potential "killer risks" is lost. Much better to *begin and end* with a focus on the big risks and leave the rest to middle management.

Final Word – Does History Repeat Itself?

This is a relevant question for boards of financial institutions. We think the answer, helpfully, is "yes and no". That asset bubbles develop periodically, organised by clever, cynical market engineers and pursued by the ignorant, financed with easy money seems to be an established fact. So in that sense history *does*, depressingly repeat itself and people never seem to learn the lessons. Indeed the biggest lesson which emerges from the study of past economic and banking crises seems to be that "the only thing we learn from history is that people never learn from history". But, and it is a big but, although *in general terms* history repeats itself, the precise nature of the build up to the crisis and its aftermath is never precisely the same, and the devil, as they say, is in the detail. Therefore, sadly, one is not protected by being fully prepared for *another crisis like the last one.* There *will* with certainty *be* another crisis but it will not be exactly the same and all the preparations may count for little. The way the Global Financial Crisis is playing out is similar but in many ways very different to previous global crises (the sheer size of the deficits, the existence of the Eurozone, changed demographics, globalisation, the rapid Asian expansion and so on). This is why boards *must* learn to anticipate conditions even if, in fact, especially if they have never happened before. Anything less is not risk management, it's living in hope.

This article is the first in a series covering risk management and the board. Future articles will cover boardroom procedures and examine specific risks in banks and insurance companies. *David Bobker is Head of Market and Operational Risk at the Asian Institute of Finance and may be contacted at dbobker@aif.org.my*