Comments submitted by Jeremy ODonnell of Risk Appetite Limited

riskapp@live.co.uk

7 September 2012

Background

Following the crisis in the financial markets between 2008 and 2009, regulators across the world were severely criticised for over-reliance on models, in particular what was seen as the flawed VaR model. While the issues that led to the crisis were much more complex than this, nevertheless it was clear that banks were not holding enough market risk capital, leading to a number of taxpayer-funded bailouts and support schemes being needed to keep the banking industry functional.

In response, the Basel Committee quickly published recommendations for three new measures: the Incremental Risk Charge, the Comprehensive Risk Measure and Stressed VaR. Together, these measures form part of what has become known as Basel 2.5, an incremental update to Basel 2 for market risk. Since January 2012, banks in Europe benefiting from VaR model approval must also apply these measures when calculating the regulatory capital requirement.

Many industry participants have raised objections to the new capital regime, claiming that there is double and, for products in scope of IRC or CRM, triple counting of risk. Some observe that the less sophisticated "CAD1" treatment available in the UK now demands less capital, so there is no incentive to adopt the more risk responsive CAD2/3 regime. The Basel Committee response is the Trading Book Review, a strategic review of the regulations that capitalise risks on the trading book. At the time of writing, the Trading Book Review is in a consultation period, but it is already clear that it will significantly alter the risk management process.

This document is written in response to the Basel Committee consultation period that ends on 7 September 2012.

Trading Book and Banking Book

The trading book is a regulatory concept that prefaces different capital treatments for risk. Broadly the idea is that if a bank enters into a position with a client with the intent to take or actively hedge the market risk, then that position is said to be held with the intent of trading and the market risk must be capitalised alongside the credit risk. On the other hand, if the position is entered into in the full expectation that it will be held until maturity without an active hedging process, then the principal risk is the risk that the client will default on their obligations under the contract, i.e. credit risk, and only this risk that must be capitalised.

There are currently no standards to determine whether a given position should be in the trading book or the banking book. Consequently firms are motivated to classify the position according to the most beneficial capital treatment. While banking book positions are outside the scope of market risk charges, they may also suffer from more punitive capital treatments of credit risk exposures. During the crisis, loss-making positions held in the trading book that were dominating the market risk capitalisation were moved across to the banking book. While it is arguable that it was no longer possible to trade the positions since the market had become illiquid, nevertheless there is a concern about possible regulatory arbitrage.

The trading book review suggests two standards by which trading book classification may be justified. Firstly, positions may be held in the trading book if the firm:

- Marks the position to market daily;
- is able to offer evidence that the price process underlying the position is liquid and observable;
- has a written policy that states what products and activities my be considered to be within the trading book
- subjects the trading book to regular internal audit review against the standards contained within the policy
- Sets and enforces limits on the position exposure
- Calibrates and reviews a liquidity horizon for the position
- Has prior experience of trading such positions or a business plan to enter a new market

The second approach is to align the regulatory treatment with the accounting treatment. Broadly the proposal is to define the trading book perimeter as surrounding all those positions that are fair valued (i.e. marled to market), although more precisely the proposal is to include all those positions for which a change in value can lead to a reduction in the Basel 3 capital resources. Some exception may be made for trades held in the banking book solely to hedge interest rate and FX exposures.

1. Which boundary option do you believe would best address the weaknesses identified with the current boundary, whilst meeting the Committee's objectives?

Of these two approaches, the former is on the surface more aligned with the existing trading intent approach, but carries with it a significant reporting burden for regulated firms and a risk that positions currently held with trading intent may fail one or more of the tests. These are largely the same tests that are later applied at desk level to determine whether a modelled approach is appropriate or if the desk should be capitalised using standard rules. However, banks will still care whether a position is on the trading book under standard rules or on the banking book.

In fact, banks may find that by adopting the fair value guideline they retain more control of the boundary. Accounting treatment is outside the domain of the financial regulator, being a matter for the auditors and international accounting standards bodies. In effect the boundary becomes a matter for the firm to decide with its auditors rather than its supervisor – and banks may prefer it this way. It is not clear whether this aligns with the committee's objectives.

Overall framework for Market Risk capital calculation

The current risk capital calculation for market risk suffers from double counting and inconsistency of approach. The trading book review proposes a three-step review and capitalisation framework that brings together existing measures in an overall coherent treatment of modellable and non-modellable risks.

Step 1 is the current model approval process to assess the overall controls in the business area for which capital is to be calculated. Until the requirements of the approval have been met, firms will have to use a newly-proposed risk sensitive standard rules measure.

Step 2 is a finer-grain review of backtesting performance and P&L attribution at desk level. This is a new review level that gives the supervisor the ability to switch on individual desks for model approval. Currently, firms may be self-selecting desks to follow the standard rules approach by excluding them from the scope of VaR Model approval applications, but under the proposals even a model approval that has been given can effectively be revoked subsequently if backtesting results are poor.

Step 3 is a risk factor level review of the risks on each desk. Provided that these risk factors can be observed at an acceptable frequency over a period of time and that the revaluation models used in the risk engines support those risk factors as inputs, then the capital can be calculated using a simulation-based risk measure. If the data is not there or the risk factor is not supported in the VaR model, then capitalisation is via an add-on calculated using a stress scenario. A special case is the default and migration risk in the portfolio, for which a separate charge (or charges) may be applied. This is where the existing IRC charge may fit into the framework (and also, if still desired, CRM).

4. What are commenters' views on the Committee's proposed desk-level approach to achieve a more granular model approval process, including the implementation of this approach for banking book risk positions? Are there alternative classifications that might deliver the same objective?

The requirement to produce backtesting results and P&L attribution for each desk may be a significant change for some firms, however this change is good as it will lead to better quality risk data for the risk manager. Desk level backtesting means that imperfections in the risk and P&L controls are more likely to be picked up. There is an investment to realise this goal however; systems must be implemented and interfaced, and the management information generated must be reviewed and actions taken. This represents a cost for the regulator as well as the firm, since backtesting results are currently subject to supervisory review and firms may get more active in challenging exceptions. Backtesting exceptions get

more frequent at lower levels in the hierarchy, as the portfolio effects that help VaR hang together no longer dominate and the specific characteristics of each business area drive the overall performance. In general firms can expect that part of their portfolio currently included in the scope of model approval will no longer be included subsequent to the implementation of this framework, and possibly that such exclusions may be significant. This doesn't necessarily mean that the business is inadequately capitalised at the moment or even in a period of stress, just that the requirement to evidence model quality will be much more stringent and consequently a greater part of each firm's portfolio will fall outside the scope of model approval than is the case today.

The committee also proposes to validate models by calculating a "risk-theoretical" P&L, a calculation of the P&L using positions at T, market data as of T and T+1 and the front office revaluation models. The risk-theoretical P&L can be compared with the actual finance P&L for the same scope of positions. Given a series of paired observations, it will be possible to analyse whether or not the VaR model is adequately capturing the P&L drivers by examining the residuals when one series is regressed on the other. More specifically, the committee proposes to analyse the following regression measures:

- The mean of the difference between the theoretical and actual P&L (unexplained P&L) divided by the standard deviation of the actual P&L; and
- The variance of the unexplained P&L divided by the variance of the actual P&L.

The committee plans to force the Finance P&L explanation process below the currently required clean and hypothetical P&L measures, to risk factor level. This gives useful feedback to the firm on the materiality of any risk factors that are missing from the VaR. The approach may require heavy investment in firm's finance functions to ensure that explanations are complete and any errors removed. P&L explanation is a discipline relatively free of standards currently and it is often difficult to separate out the explain contributions in the presence of cross risks. For the process to be meaningful, firms should use VaR model parameterisations of curves in the front office models – something that may require quant development.

The proposal does not require that the risk-theoretical P&L is calculated using the firm's historical simulation engine (a VaR model P&L). Such a process is more or less equivalent to running a historical simulation with one scenario, with the additional feature that the market data is not historical relative to the as of date for the positions. This analysis can

provide a more complete view on whether the P&L drivers are being adequately modelled in the VaR model. It serves the purpose of testing for capture of material P&L drivers as above, while additionally testing that the VaR model is able to capture the price behaviour of the portfolio given a set of risk factor movements – another important test not necessarily covered by backtesting. This will not require any additional model development work and does not rely on potentially data intensive cleaning processes that may be difficult to validate. For these reasons this procedure is recommended as an alternative to the committee proposal.

The definition of desk for the purposes of the framework will be aligned with internal organisational boundaries. If firms do not already recognise the importance of a common and consistent hierarchy between the front office, product control and risk management then it will become self-evident when applying the framework. It will be necessary to report risk and clean P&L for a common scope of positions at an aggregation level of desk that relates to front office organisational boundaries. The only sensible way to approach this is with a common hierarchy. To attempt to do this without a common hierarchy process is to invite failure (i.e. standard rules across a large swathe of books and a poor regulatory relationship).

The consultation recognises that the requirement to calculate market risk capital for FX and commodity trades on the banking book does not fit easily into the proposed framework and proposes three alternative approaches:

1. assumed model approval
2. standard rules charges
3. assessment at asset class level for FX and commodities on the banking book

Option 3 might be the best of these proposals, but it suffers from the weakness of being too granular a decision point. If there is one trade that does not fit well with the model, then the whole asset class is excluded. The optimal approach is to allow banks to blend options 2 and 3, so that troublesome trades can be moved outside what we might term the modelled perimeter and have standard rules charges applied. This is in fact the approach that is currently available, with the additional control of more rigorous assessment of the modelled perimeter.

Risk Measure for Capital Purposes

The next problem dealt with by the trading book review is the disconnect between actual losses and VaR. The 10-day loss can be expected to exceed the 99% confidence level 10-day VaR measure between 0 and 2[1] times every year (using non-overlapping observations). The standard multiplier of 3 was intended to allow for losses to exceed the VaR without exhausting capital. However those involved in drawing up the regulations have been quoted as saying that there was never any science behind the selection of 3 versus any other number.

The Basel committee now proposes to use an alternative risk measure, the Expected Shortfall or ES. The ES is the average of the observed losses in the tail of the P&L distribution, where the tail begins at the alpha-percentile.

The value of alpha to be used has not yet been proposed, however obvious candidates are 0.95 and 0.99. There is no proposal as yet for the multiplier, which is likely to be calibrated to achieve whatever capital result is desirable at industry level (either flat or an increase of a desired level).

8. What are the likely operational constraints with moving from VaR to ES, including any challenges in delivering robust backtesting, and how might these be best overcome?

The idea is that the ES is a better estimate of the capital requirement as it reflects the extent of the tail losses. It is also what has been labelled in the academic press a coherent risk measure. However there are some problems with this approach.

- It will still be possible for losses to exceed the ES
- For values of alpha less than 0.99, it will be mathematically possible for the ES to be less that the 99th percentile, leading to a reduction in the capital requirement
- For values of alpha exceeding 0.99, the ES may be unstable, making it hard to build credibility for the risk measure and integrate it with risk-based limit setting
- Revaluation frameworks tend to become less accurate for more extreme scenarios

[1] 25 Binomial trials with probability of success = 0.01 at the 99% confidence level

It seems a tremendous risk to adopt a different measure as the basis of the capital calculation without any understanding of how that will integrate with the financial services target operating model. Attractive though it might be in the academic literature, unlike VaR there has been no widespread adoption of ES across the industry in advance of its adoption into the regulatory framework. If the credibility of the measure cannot be established, it will create a permanent disconnect between the capital calculation and the risk management process. This is evident also in the requirements for backtesting. If the ES calibration is such that it exceeds the 99% VaR, then exceptions will be less frequent and the backtesting process will cease to be effective as a monitor of capital adequacy efficacy.

It is noted that despite the fundamental questions this raises, the committee is not requesting comments on whether the change is the only available solution to the undercapitalisation of tail events when using VaR.

Another knock-on effect of ES is that accurate valuations in the tail will be demanded. The Basel committee expects that every implementation will use full revaluation across the whole portfolio, continuing the regulatory obsession with full revaluation in place of scenario matrix or sensitivity-based valuation methods. This will be costly for many firms and ignores the benefit that comes from being able to conduct what-if analysis within a sensitivity or scenario matrix framework. Variance-covariance methods will no longer qualify for regulatory approval.

The main benefit of adopting ES as an alternate to VaR is that it is a risk metric that can easily be calculated from existing VaR pnl vectors that have been produced in support of a simulation-based methodology. However, this does not mean that the implementation can be completed at zero cost, since many reporting suites and end user solutions will have the current VaR metric hard wired into them. Additionally, the VaR process will have to change almost beyond recognition to support the other proposals in the consultation, so there is little real benefit here.

An alternative approach to reflect the tail of the distribution may be to increase the minimum time series window from 1 year to something more representative of the risk, perhaps four or six years. Non-overlapping scenario observations would limit the impact

of this change on the pricing engines (statistically, overlapping observations are in any case biased[2]).

One of the reasons that losses may exceed the ES is if the ES is calculated using a time series window that is relatively benign, as was the case in the period 2002-2006 and may happen again. The Basel Committee has considered this weakness and determined that the ES must be calculated using a time series window that represents a stressed period for the portfolio. This offers a solution to this problem but throws up a few of its own:

Stressed VaR is calculated using just 250 days of data and is therefore quite volatile with a large standard error. However the impact of the stress is reduced if a larger time series window is used, unless the enlarged window includes another period of market stress. For instance, if the number of scenarios is doubled but none of the new scenarios are in the tail, the 99% VaR measure goes from being the average of the ranked second and third worst of the original period to the ranked fifth worst. So, there is no easy way to stabilise the distribution without diluting the VaR measure.

The number of scenarios also makes the move to the ES measure less of a quantum leap than it may appear. With 250 observations, each observation has a probability density of 0.004. The worst ranked observation is the 99.6 "percentile", or 996th permille[3], the second worst 992nd permille, third worst 988th permille. An estimate of the 99th percentile may be obtained by averaging the ranked second and third worst observations. The ES calculation at an alpha value of 0.988 is the average of the ranked first, second and third observations, i.e. just adding the worst loss to the existing calculation. A lower alpha value, for example 0.98, would extend the ES calculation to include the fourth and fifth worst losses. This may lead to a more stable ES and different metric, however adding weight to the lower ranked losses may have the effect of reducing the ES below the current 99% VaR, if the gap between the worst loss and the current 99% estimate is smaller than the gap between the 99% and the midpoint between the fourth and fifth ranked losses.

---

2 for n-day observations of order i=1…N, the first/last (n-1)-ordered daily movements contribute to the computation of i/N-i of the overlapping observations, whereas all the intermediary daily movements contribute to the computation of exactly n of the overlapping observations. Hence the first n-1 and final n-1 ordered observations carry less weight in the final calculation and the resultant VaR computation is biased

3 The equivalent of percentiles when dividing the cumulative distribution into thousandths instead of hundredths

As mentioned earlier, firms and the regulators will have to consider the impact of migrating from VaR to ES on the backtesting programme. Every firm with model approval currently reports backtesting results at consolidated level and one level below (usually an organisational or asset class division). Many also conduct backtests at lower levels, and some do this for multiple regulated entities. For each regulated entity, the back test results determine the multiplier to use on the VaR. The multiplier is used by the regulator to reflect the fact that tail risks are not fully captured by VaR, and also to secure additional capital when the model is seen through backtesting results to be performing badly, and to otherwise correct for what are seen as process control gaps and to incentivise firms to close them. However the proposed regime is somewhat different and envisages a more black and white decision. Backtesting will be conducted at desk level. Either the desk level backtesting will support the inclusion of that desk in the entity portfolio, or it won't. If it doesn't, there is no provision to include the modelled number using a multiplier – the risks on the desk will be reflected using standard rules.

The committee is still trying to determine what the criteria for passing the backtest may be. One option is to continue to back test with VaR. The statistics of testing with expected shortfall are more complex because the probability density beyond the expected shortfall is not known. The committee still intends to reflect backtesting performance via a multiplier, so it seems that there will need to be an amber zone for less good quality backtesting before the red zone requires standard rules to be applied. One approach would be to change the percentage to be used when calculating the standard rules floor. For instance, if capital is floored at 60% of the standard rules charge, desks in the amber zone may require 80% of the standard rules charge as a floor, easing the transition to the red zone (when the percentage becomes 100%). The committee is also considering ways to reflect the size of backtesting exceptions in the capital requirement. This is an interesting idea but may be flawed if it does not also consider the size of 1% exceptions. In stable markets, unusually large losses are often followed by almost equal profits the following day as the market corrects. This scenario is obviously a lot more benign than repeated sizeable losses in stressed markets without any materially offsetting gains. In the latter case, it seems sensible to define an alternate trigger for entering the red zone than the number of exceptions, based on cumulative losses over the liquidity period exceeding or likely to exceed the capital requirement for that period. In this event the whole desk could switch over to standard rules. The criteria for exiting standard rules would need to be designed to reflect

any remediation carried out by the firm, for instance a change of the stressed calibration to include the loss period, or a development of the VaR model to include the P&L driver that was missing.

The cost of implementing a migration to ES is likely to be significant for most firms, as it goes way beyond calculating a different statistic on a pnl vector. Software features that highlight given percentiles will need enhancing to show the ES calculation; end user reports and core systems functionality that hard-wires the VaR metric will need to be changed to show both VaR and ES or just the ES. Backtesting programmes that do not operate beyond the current regulatory minimum requirement will need to be expanded significantly. In doing so, firms may discover that the hierarchy solutions they have been using are suboptimal and replacing them takes significant effort across the firm.

Calibration

When calibrating the ES using a period of stress, the regulator is biasing the measure by predetermining that the "random" scenarios must be drawn from a period of market stress. Currently this is defined as a continuous period of one year for which the organisation can demonstrate that the markets driving the risk of the portfolio in question was at stress levels. At the time of Basel 2.5, there had been a long period of market stability stretching from early 2002 into 2007, and a more volatile period running from August 2007 through to 2009. In this context it may have appeared sensible to require firms to choose the most stressed 1-year period out of the roughly two years. Unfortunately the markets are still stressed, and not showing any signs of getting less stressed, as the indirect repercussions of the original shocks in 2007 and 2008 continue to reverberate around the market and manifest themselves as sovereign debt crises (multiple), potentially a Eurozone break-up crisis, and so on. It is therefore much harder to pick the continuous year that represents the period in which the markets were most stressed. In fact, although it has been the convention for many years now, there is no reason why the sample period for a historic simulation has to be continuous. An organisation should be able to pick several periods of market stress, in the same way it picks stress test scenarios, and in this way build up a larger body of historic scenarios for use in the risk calculation. If each scenario period chosen was 125 days (6 months) long for instance, regulated firms could pick 4 scenario periods to produce a risk metric based on 500 days of data.

This does leave the problem of how to identify periods of stress, but this is a problem now as well. Firms are expected to do analysis back to January 2007[4] to demonstrate that their stressed period is appropriate. The trading book review proposes explicit analysis be conducted, although some attempt has been made to reduce the dimension of the problem. Instead of considering the full range of risk factors in the portfolio, firms may consider a reduced set of risk factors that represent the key risks in the portfolio. This analysis needs only minor changes to identify four periods of 125 days instead of one period of 250 days. It may also be easier to persuade the regulator of the appropriateness of shorter periods using qualitative arguments, although there is no scope for this flexibility within the current consultation draft.

Another challenge for the ES method is that the scenarios calculated for the stressed period may require a significant degree of proxying. Regulators do not like proxying as it understates the volatility of the scenarios. Firms do not like proxying because it understates diversification benefits. Either way, the more proxying there is, the larger the error in the tail scenarios. The consultation paper proposes an alternative *Indirect* method to calculate the ES for a stressed period. The approach requires the calculation of ES for the stressed period only for the reduced set of risk factors (and by implication a high degree of proxying). A loss scenario from this period – perhaps the worst loss, but it is not clear - is then scaled using the ratio of ES for the most recent 250 days using all risk factors to the ES calculated using the reduced set of risk factors. Clearly this may increase or decrease the result, depending on whether volatility of real risk factors or diversification benefits has the larger overall impact.

5. What are commenters' views on the merits of the "direct" and "indirect" approaches to deliver the Committee's objectives of calibrating the framework to a period of significant financial stress?

The practicality and effectiveness of such an approach is open to debate. Firstly and fairly obviously, there is no reason in particular that using this scaling approach will produce an answer that is any more correct than an ES calculated for the identified stress period using all risk factors (this is known as the *direct* approach). The idea behind the indirect approach

[4] Prior to 2007, you have to go back to 2001/2 to find a stressed market (Iraq war, 9/11, Dotcom bust). By the time you are that far back, there is a strong argument that the market was structurally different – notionals outstanding in the credit derivatives market were a lot lower, for instance – and historic simulation techniques are not appropriate)

is that an accurate ES can be calculated over a recent data set. The implicit assumption is that the relationship between ES on reduced risk factors and ES on the full set of risk factors is stable over time. This should be a safe assumption if the additional risk factors are not cumulatively material in the portfolio, however if that was really the case there shouldn't be any need to model them at all – a fixed scaling factor could be determined based on periodic calibration. Moreover, the whole idea of calibrating to a stressed period is that the current market conditions cannot be relied upon to accurately capture the potential loss during market stresses.

Given the questionable legitimacy of the approach, the disadvantages seem to be significant. Searching for a stressed period on the basis of reduced risk factors is not the same as searching on the basis of first order sensitivities. Full revaluation calculations are required on the reduced set of risk factors for the entire search area. If you take 1 Jan 2007 as the starting point, by the end of 2012 1500 full revaluations are required across the portfolio. Once you have selected the stress period, you then need an additional 250 full revaluations using this time the complete set of risk factors, and another 250 full revaluations using the complete set of risk factors for a 1-year standard VaR window (this is incremental unless the stressed calibration period overlaps with the most recent 250 days). All in all, depending on the size of window used for the current standard 1-day VaR calculation, the degree of overlap between the stressed period and the most recent 250 days, and the frequency with which firms will be required to re-evaluate their stressed period, this could quadruple the amount of risk engine work required (and counting) for firms using simulation methods. Firms that lack complexity and consequently have a model approval to use the variance-covariance approach would face an even greater step change, and it is likely that such firms will need to be allowed to revert to the new risk-sensitive standard rules.

While it is true that the approach could drastically reduce the amount of time series data that needs to be captured across the whole historic period, a mapping mechanism must be maintained to relate the full set of risk factors to the reduced set. For instance, if the full set of interest rate risk factors is mapped to the first principal component, the risk engine scenarios must recognise the relationship between scenarios in the principal component and scenarios in the atomic risk factor set used for revaluation.

The most worrying aspect of the indirect method is the impact that it will have on the transparency of the capital calculation. Currently, the capital calculation is aligned with internal risk management practice. With the right infrastructure decisions, a firm can hope to explain the changes in capital requirement in terms of changes in the portfolio and market data. This connection becomes more difficult with the direct ES method, since the capital calculation will require an understanding of each of the scenario outcomes in the tail. With the indirect method, it is possible that it will no longer be possible to explain the capital charge.

If there is regulatory concern that the search for the stressed period is not comprehensive enough, it would be better to allow the search for the stressed calibration period to compromise on accuracy rather than or as well as scope. For instance, if the calculation could be based on a delta-gamma or scenario matrix approach for the search, this could be completed on a dedicated risk platform at a fraction of the cost of supporting 2000+ daily full revaluations. Once the period has been identified, the direct approach can be used to derive the ES and therefore the capital requirement.

Market Liquidity

A significant driver of loss in the credit crisis was the fact that positions that had been trading freely became illiquid. The issue of liquidity is acknowledged in the current VaR calculation by using a horizon of ten days. During the market crisis, this was found to be woefully inadequate, as markets became one-way. Consequently, the committee wishes to capture the concept of liquidity more explicitly.

The committee proposes to capitalise two separate liquidity effects, of which only one is considered in these comments. The variation in liquidity horizons in different markets will be captured using a set of five liquidity buckets ranging between ten days and one year. These will allow for different levels of liquidity in different markets, similar to the product-based liquidity assessment in the IRC calculation. However, rather than determining the liquidity based on a broad product classification, the liquidity horizon will be based on the bank's internal assessment of the risk factor liquidity, subject to regulatory floors at asset class level. Firms will also have to account for the size of positions held and the additional time it will take to unwind large positions without impacting market prices.

The committee has not yet determined how to aggregate results from different liquidity horizons, but appears to favour one of the following approaches:

- simulation or historical observations of risk factors at the different horizons through to 1 year (1)
- observation of the risk factors at different scaled horizons with subsequent rescaling to the correct horizon (2a); or observation of the risk factors at a single horizon with scaling (2b)
- observation of all risk factors to a single weighted average liquidity horizon (WALH), where the weighting represents the proportion of risk at each horizon (3)

2. What are commenters' views on the likely operational constraints with the Committee's proposed approach to capturing market liquidity risk and how might these be best overcome?

The committee acknowledges that collecting sufficient quantities of unbiased historical simulation data at long time horizons will be problematic (bearing in mind that this must also be calibrated to a period of market stress). Even when making observations at a scaled horizon, 250 non-overlapping observations at a 1 year horizon, scaled down to 25 days, will require 25 years of valid time series data. Equally, scaling techniques that are materially effective at a 10 day horizon are unlikely to be correct for 1 year. For method (3), finding a metric that doesn't represent a fudge seems problematic, since this would strictly need to reflect the absent data that has driven us down this route in the first place. With the exception of method (1), each of the methods is either infeasible or dislocated from the underlying risks and market data to some degree.

Option (1) is more interesting, since this simulation procedure, although onerous compared with prior market risk implementations, is a fundamental component when modelling credit risk exposures. The differences are firstly that a credit risk simulation usually makes do with a reduced set of risk factors appropriate to the scope of positions (usually derivatives) that are being simulated, and secondly that the credit risk system will typically use custom models optimised for speed in place of the front office models preferred in market risk systems. Nonetheless there would seem to be a real opportunity to merge the credit risk and market risk engines in fulfilment of this new regulatory requirement. The market risk portfolio value can be modelled through its life and losses accumulated in each liquidity bucket through to the 1 year liquidity horizon. The ES can be observed from the ranked accumulated losses.

While risk is assumed to roll off once it reaches its liquidity horizon, the committee will allow adjustments to ensure that hedges stay on for the same time as the positions they hedge, through what they term an "instantaneous shock" method. This requirement assumes that a hedge may have a different liquidity horizon to a trade. Since the proposal is to set liquidity horizons by risk factor rather than product, it would seem that hedges with the same risk factor as the trades they hedge should have the same liquidity; if it does not have the same risk factor, it is not a valid hedge from a liquidity perspective and for a stressed calibration set it is correct to let the hedge roll off at its liquidity horizon leaving the position exposed. It may be possible to model a more sophisticated approach than this and the industry should be encouraged to do so.

The advantage of a simulation is that prices are able to follow a natural evolution using arbitrage-free processes that are also used by the trading function for pricing. No arbitrary scaling assumptions are required, and sampling optimisation techniques can be used to focus attention on the tail of the distribution. The question of aggregation of different liquidity horizons is neatly dealt with by summing losses on each scenario path. In short, this is the only approach that retains transparency when adjusting the ES measure for liquidity.

Diversification Effects

One of the underlying principles of VaR and a key foundation of its risk sensitivity compared to standard rules measures is its ability to capture the optimisation of risk and return and recognise the benefits of diversification across asset classes. However there is a concern that calibrating VaR calculations using historical data may imply benefits that are not retained during periods of stress, as correlations polarise. The committee now proposes to restrict diversity benefits so that they cannot operate indiscriminately across asset class boundaries. Instead, the regulator will set asset class correlations to be used when aggregating across asset classes.

There seem to be three problems with this approach. Firstly, there is no guarantee that aggregation using asset-class correlations will result in a larger capital number than a fully diversified ES observed in the tail of the full portfolio loss distribution. To guard against this circumstance, firms must also calculate this number and the capital required will be the maximum of this number and the aggregated asset class numbers. Secondly, the asset-class correlations may disallow specific hedges that are effective in all expected market conditions, such as BP shares versus oil futures. Thirdly, the classification of the portfolio by asset class may not be straightforward – hybrid products exist, and most products carry interest rate and FX risk. The consequent treatment is not trivial.

The two options proposed by the regulator are:

- Shock all the risk factors belonging to a particular asset class simultaneously or obtain the partial derivatives of the risk with respect to each asset class.
- Assign each position to a primary asset class that drives the material price behaviour

6. What are commenters' views on the merits of the desk-based and risk-factor-based aggregation mechanisms to deliver the Committee's objectives of constraining diversification benefits?

In the first case, the reference to partial derivatives is puzzling in the context of the prior statement about use of full revaluation, since partial derivatives are necessarily an approximation. This method is also subject to the problem that cross-asset cross risks will not be captured in the model, and the risk classification is not aligned to the business organisation.

In the second case, risk classification can be aligned to business organisation and system divisions but aggregation using the calibrated correlation parameters may have less meaning. Given that the model approval perimeter itself will now be set along granular organisational boundaries, it may be that this method is more transparent in the new framework.

This part of the new modelling framework is probably the most dissatisfactory. While the objective may be defensible, the treatment is crude and not at all risk sensitive. Consequently the alignment between top of the house capital requirements and risk reporting is lost, and firms are forced to partial failure of the use test, which currently requires firms to demonstrate that the methodology used for capital reporting is also used for internal management of risk positions. Most distressing of all, the breakage is at the cross-asset level where senior managers need to take a view on combining risks that may make sense in their asset class silos but not in combination at a group level.

As an alternative, it would be better to do directly what the regulator aims to do by setting calibrated correlations for cross-asset class aggregation: run simulations under different correlation parameterisations and see what the worst case scenario can be. The correlation parameterisation becomes another degree of freedom, with a random draw at the beginning of a simulation dictating the parameterisation to use. Each path is weighted equally regardless of which parameterisation has been used and the results can be ranked across all parameterisations. The aggregated results at cross-asset class level are obtained using the simple vector sum as they are today, transparency is retained and the use test is intact.

Credit risk

Recent improvements to the market risk framework have focused on the growth of credit derivatives trading over the last decade. First there was IDRC (a default risk charge) then IRC (a charge that also covered transition risk) and then CRM, a comprehensive risk measure for correlation trading products (known as APR, All Price Risk, in Europe). All of these charges aimed to better capitalise the traded credit risks in the trading book. Under the new framework, the committee faces a choice as to whether to continue to allow these charges or to declare that these risks have to be capitalised using the new risk-sensitive standard rules. In some ways it will make the implementation and maintenance of the

model more straightforward if risk-sensitive standard rules charges amount to a charge roughly similar to the existing CRM floor. However since the direction of the whole market risk process will be to make greater use of simulation methods, it makes sense to retain the experience and discipline of IRC and CRM simulations and keep these risk factors in scope for the simulation process. If the risk factors subsequently fail the control tests in the framework then the appropriate path can be followed, as with any other risk factor.

The committee cites a number of disadvantages of this approach. Firstly, default and migration are discrete, low-frequency events that are difficult to backtest – i.e. a backtest can show a "false positive" (Type 1 error) that implies the model is performing correctly when it is not. Secondly, the percentile used for the ES measure will have to be set at a high level if these risks are included. Secondly, the process for correlating credit event risk factors with market risk factors will naturally favour companies with a long history and underplay the role of failed companies (because there is no current data for failed companies). Thirdly, the model becomes more difficult to supervise if an integrated approach is used.

The important thing is that the incentive for the risk management function to accurately simulate the risk in the portfolio is retained. Yes, the supervision of an integrated model is harder, but the P&L function is integrated, so the risk treatment must be also. Yes, correlated data will be biased so the industry must develop a treatment that corrects for it. And market risk factors also include event risk that is being captured now that we have an ES measure, so credit events can be modelled consistently without necessarily moving the percentile level (it will have to be high anyway to prevent an overall reduction in risk capital).

It would in fact seem a little perverse to migrate the rest of the market risk capital approach to what is effectively an evolution of the IRC and CRM frameworks and not include these charges in scope.

Process and Systems implications

The process and systems implications of the trading book review are wide-ranging and affect many areas of banking activity. Firms will need to plan well in advance to meet the

needs of the trading book review. Many times, firms rush to bend their existing systems and processes to meet the needs of the latest regulatory hurdle. This approach may well be less costly in the short term, but it will push firms down the route of adopting the liquidity treatments that can be bolted onto the existing VaR systems, rather than meeting the challenge of integrating credit risk and market risk platforms. This will lead to an inferior capital measure and compromise the risk management process. This is a fundamental review of the trading book and firms should expect fundamental changes to their systems infrastructure and processes.

Summary

The Trading Book Review was undoubtedly needed after fifteen years of capital calculation using VaR models, regardless of the issues encountered during the crisis. The emphasis on controls, P&L explanation and VaR model validation are certainly welcome. However the current proposals contain several questionable features:

- The proposed migration to ES apparently discounts the fact that this measure has been largely ignored by the industry
- Few of the proposals for liquidity adjustment are elegant, transparent or implementable
- The proposals to limit diversification benefits through regulator-set correlations disregard the importance of a risk-sensitive measure in the use test. As a consequence, traders will be faced with a choice of hedging their risk or hedging their capital requirement
- The requirement to use full revaluation for all instruments will place an undue burden on the hardware requirements for banks with any kind of scale
- The reliance on Finance P&L cleaning by risk factor will add pressure to what is typically already an overloaded support function in the firm

Instead, the committee should

- Promote integration of market risk and credit risk simulation frameworks for integrated risk management
- Consult with the industry on appropriate risk measures in such frameworks
- Embrace diversified risk measures supplemented with correlation scenarios
- Allow wider use of scenario matrix approaches than is encouraged today
- Require VaR Model validation using the risk engine framework

Whatever the final regulations are, most firms will have a significant programme of work to meet them. There is a real danger that firms will fail to react now and will then be surprised by the scale of the change when the committee issues final proposals. In part this will depend on the elapsed time between the final proposals and the implementation date, but in any case this is a period of uncertainty for most firms as regional and local regulators all work to finalise their enactments of the Basel regulations. The committee commented on the Basel 3 implementation that it was difficult to assess whether the target of a level playing field was being achieved as multiple drafts of regulations circulated in each jurisdiction. For this reason, many firms leave it until the implementation year to scope out the implementation. This would be a mistake and lead to rushed implementations and poor choices, with firms favouring approaches closest to their current VaR implementations rather than making strategic choices that optimise the use of resource and the benefit to the organisation. The trading book review should be thought of as getting a new model approval. It is a multi-year project that requires considerable resource and cross-departmental prioritisation. Firms need to start moving themselves to an interim state that will ease the path to meeting the eventual regulations. Those that don't risk being left with inefficient capital usage and broken risk management processes. It is hoped that national supervisors will recognise this and drive early definition of project work to support the review.

About the commenter

Jeremy ODonnell (London, UK) is the proprietor of Risk Appetite Ltd, a provider of consultancy services in the financial industry. Jeremy formerly worked in technology, finance and risk management roles at Barclays Capital, Merrill Lynch and Gulf International Bank. A specialist in risk systems project management, he has overseen the implementation of market risk systems at all of the preceding financial institutions. His current focus is on the implications of the Basel process for risk-based capital in large banks, with particular reference to system architecture and process issues.

