

7th September 2012

Mr Wayne Byres
Secretary General of the Basel Committee
Bank for International Settlements
Centralbahnplatz 2
CH-4002 Basel
Switzerland

RBS Risk Management
280 Bishopsgate
London EC2M 4RB
Switchboard: 020 7672 1239
www.rbs.com

Via email: baselcommittee@bis.org

Dear Mr Byres

Consultative Document: Fundamental review of the trading book

It is a pleasure to enclose The Royal Bank of Scotland's ("RBS") response to the first consultation on the Fundamental review of the trading book ("Fundamental review"). The review is necessary to strengthen the robustness of the financial system and we welcome the opportunity to comment.

RBS has also contributed to the Joint Industry response (by AFME, IIF, ISDA, and TCH) and the responses by the BBA and the EBF, which we endorse.

From our attached response I would like to highlight the following comments:

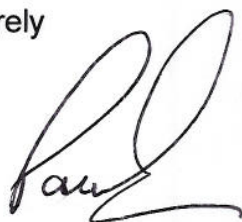
- RBS would have preferred a more fundamental discussion about the role of the trading book/banking book boundary and possibly more innovative ways to overcome the cliff effects that a binary boundary brings. Should the binary boundary remain we feel that the evidence based boundary is more likely to achieve the Committee's objectives.
- RBS agrees with the proposal to factor market liquidity into Value at Risk ("VaR") or Expected Shortfall ("ES"). Whilst a variable liquidity horizon framework is more consistent with how we manage risks internally, we do not feel that other approaches should be dismissed at this stage. We would support a closer engagement with industry on whether simpler approaches might be appropriate from a financial stability perspective.
- RBS supports the simplification of VaR and Stressed VaR into one model.
- RBS supports the move to ES. In our response, we point out some of the issues that regulators should consider when implementing this approach. One key issue

is the move to full revaluation. While we support this move in principle, we believe that it would bring significant implementation and cost issues for many banks. We suggest that the option to use sensitivity based approaches should be allowed where their effectiveness can be demonstrated.

- RBS is uncomfortable with the proposals on diversification. Our preferred option would be to maintain modelling of diversification. A sliding scale approach would be a second best alternative.
- On credit risk, RBS feels that the door should remain open for firms to determine whether an integrated or separate modelling approach is the better option.
- RBS is supportive of the model oversight proposals, though an approach that mitigated cliff effects should be considered.
- Regarding the revised standardised approach, our current view is that the fuller risk factor approach would provide improved financial stability, as it would be less likely to provide incentives to diverge from good risk management behaviour. It would also be easier to implement for most advanced firms. We do recognise though that for some smaller firms this approach may not be appropriate, particularly in jurisdictions that do apply Basel to non-internationally active banks. Therefore RBS would additionally support the retention of something similar to the current standardised approach, which might for example be available to firms that are not subject to a Global or Domestic Systemically Important Financial Institution categorisation or some other measure to ensure that larger institutions could not use such a simplified approach.

RBS would of course be very happy to provide further details on any of these issues to the Basel Committee.

Yours sincerely



Paul Ingram
Global Head of Market and Insurance Risk
Global Market and Insurance Risk
The Royal Bank of Scotland plc

Direct tel: +44 20 7672 3354
Email: paul.r.ingram@rbs.com

RBS Response to the Fundamental Review of the Trading Book

1. General comments

In general terms we welcome the Committee's proposals for both simplifying and improving the modelling framework used to capitalise positions held in the trading book. We appreciate and share many of the concerns expressed with the current models, even if we may not always entirely agree with the suggested alternatives. In particular, we are fully supportive of the intention to rationalise the number of models used to calculate Pillar 1 market risk capital, and thereby eliminate much of the double-counting of risks that is a feature of the current framework.

RBS would urge the Basel Committee to give enough time for implementation of the final rules, which are likely to require large system changes. In particular we suggest that there is a further round of consultation of more detailed proposals prior to the intended Quantitative Impact Study ("QIS"). This would help improve the relevance of the QIS exercise. Banks are likely to need about six months preparation for meaningful results after publication of detailed rules.

Given the range of enhancements to the market risk capital modelling framework outlined in this consultation paper, we assume that there will be a re-assessment of the corresponding capital multipliers. In particular, we would note that the current minimum capital multiplier of 3 was introduced by the Basel Committee based on concerns around a number of potential weaknesses with the use of 10-day VaR models. Specifically, that:

- *The past is not always a good guide to the future;*
- *The assumptions about statistical "normality" built into some models may not be justified, i.e. there may be "fat tails" in the distribution curve;*
- *The correlations assumed in the model may prove to be incorrect;*
- *Market liquidity may become inadequate to close out positions.*

The final two points in particular appear to be explicitly addressed by the new proposals. In addition, typical VaR models currently used in the industry do not assume normality — in fact; probably the most widely used methodology is based on the observed empirical distribution of risk factor returns. This, taken together with the use of an ES measure calculated directly from the modelled loss distribution, should remove or at least significantly alleviate this concern. Finally, the use of a stressed model calibration ought to mitigate concerns around the extent to which the past can be used as a guide to the future, although of course it is not possible to entirely eliminate this inherent model weakness.

You will be well versed in the industry's views on the CVA volatility charge proposed in Basel III. We support those views and would urge the Basel Committee to revisit the CVA volatility charge in a timeframe compatible with the fundamental review.

In the following, we provide responses to each of the Committee's questions arranged by broad topic.

2. Trading book/banking book boundary

Q1. Which boundary option do you believe would best address the weaknesses identified with the current boundary, whilst meeting the Committee's objectives?

From our perspective, there are few conceptual reasons to choose between the two proposals, especially if the use test was to be de-emphasised by regulators. We have a slight preference for the evidence based approach, as explained below.

Our primary principle guiding the way we have approached this issue is that a framework where risk management view diverges from capital requirement is not desirable, assuming of course that risk management takes economic risk into account. Where accounting differs from a risk management view of the world (and therefore from an economic risk perspective) accounting would not be a good guide for boundary classification. Under such circumstances, such as split hedges, we feel that the risk management view should prevail. This could become a particular issue if accounting standards changed markedly from what they are today. We would therefore not support a complete reliance on accounting standards to determine risk management processes.

We also feel that an expansion of the trading book to incorporate many more firms and treasury activities, needs to carefully consider the additional risk of unintended consequences.

Nevertheless, we feel there would still be merit in looking at approaches that contain fewer cliff effects than the present proposals. This could for example be achieved by having:

- a trading book restricted to those products that only contain very liquid risk factors.
- a further book with a one year liquidity horizon, which somehow takes into account market risk without pushing the activities into a fully fledged trading book.

We feel such an approach would have the benefit of also being applicable to portfolios held as liquidity buffers, which clearly contain market risk, but which are also potentially held indefinitely.

We note that our views may change once more details of the new trading book proposals become available. In particular possible changes to the calculation of capital on interest rate risk in the banking book are another factor that could change the impact of the Fundamental review.

In relation to the two options being consulted on, we would make the following comments:

2.1 Trading evidence-based approach

This approach is preferred to the valuation-based approach as it has the potential to align firms' risk management processes with regulatory classification. However, it is not clear to us whether the current proposals are simply a further codification of existing best practice in line with the proposals in BCBS 116 *"The Application of Basel II to Trading Activities and the Treatment of Double Default Effects"*. We are of the view that the good risk management policies that are applied by many banks already can be used to address many of the issues raised, in particular relating to valuation, dealing mandates, limits etc. However if the intention of the Basel Committee is to go much further, then it would not be clear to us what additional benefits would be achieved. We would hope that any additional evidential requirements would be practical enough to make it easy to understand and be objective, being neither too onerous nor too bureaucratic. Also, it is unclear how much opportunity there would be to elect classification on initial recognition.

We also believe that a boundary should be permeable to some extent as strategies and markets change. On the other hand, we do not think it would be appropriate for products that were correctly put into one category at inception to switch back and forth between books based on changing trading evidence. We

believe that a permeable boundary can be supported by firms having policies in place to control the movement of assets across the boundary, with reporting requirements on assets which have been moved. This being a conceptual consultation we would question whether too much emphasis is placed on the objective of reducing possibilities of arbitrage. We say this because in future, with the introduction of liquidity horizons and incorporation of tail events, it will be far less certain that one treatment offers a lower capital treatment than another. Therefore, on balance, we feel that an alignment between internal and regulatory categorisation will be more beneficial from a financial stability perspective.

At present we cannot comment on how the trading evidence-based proposal might work for new products. With stricter requirements to evidence trading required upfront, it is difficult to see how new products could be accommodated in the trading book until trading is well established. The framework should allow for such products to be re-designated to the trading book when they meet appropriate criteria. Alternatively, the framework could allow such products to be designated to the trading book at initiation on the basis of the risk factors involved.

2.2 Valuation-based approach

We generally agree that the valuation-based approach can be seen as more objective, given it must be agreed with auditors. Also, as the accounting standards deal with financial instruments generally and not specific products, this approach would more easily accommodate new products. However, basing categorisation on a different framework with different objectives could lead to categorisation issues. In practice, difficulties similar to those encountered in the current approach would probably emerge.

We accept that arbitrage is unlikely to occur after initial recognition under the valuation-based approach. However, there is some flexibility to determine accounting treatment upon initial recognition, with the potential for arbitrage at inception. Assets could not be moved across the boundary once recognised, unless there was a change to accounting standards. We have already noted that this is not conducive to good risk management from a market risk perspective.

Furthermore the Committee should be mindful that the rules are set by accounting standard setters and monitored by auditors, not regulators. Accounting standards change over time and the objectives for accounting standard setters are not necessarily aligned with financial regulators. Moreover, accounting standards are not consistent globally and this may result in different regulatory classifications across firms. We therefore think that ultimately this approach would be less consistently applied than the trading evidence-based approach.

Finally the valuation-based approach does not align with RBS's current risk management processes. Many products that are classified as held for trading or available for sale are not managed on a short term basis. This approach could therefore significantly increase the size of the trading book, leading to large operational issues. It would also bring many more firms into having to run trading books, even though this will not be warranted by their business model and would not be appropriate from a risk management perspective. We also wonder whether the incentives provided to firms by having such portfolios under market risk rules would create less financial stability instead of more. We note Andrew Haldane's speech entitled "Accounting for bank uncertainty"¹ and feel that the valuation-based approach could exacerbate the problems that he has identified.

We also feel that the Committee would need to address the issue of hedging non fair valued assets with fair valued assets and whether to include available for sale assets in the trading book.

¹ <http://www.bis.org/review/r120123a.pdf>

2.3 Conclusion

In conclusion, we would re-iterate that our preferred solution is alignment to risk management. If the capital requirements were correctly calibrated then boundary and possible arbitrage would be less of an issue. We noted earlier that this may therefore have been an area where the Basel Committee might have benefited from looking at more ways to address the underlying concerns.

3. Factoring in market liquidity

Q2. What are commentators' views on the likely operational constraints with the Committee's proposed approach to capturing market liquidity risk and how might these be best overcome?

In high-level terms, we are aware of two apparently quite different modelling approaches to capture the first-order effects of market liquidity:

1. Variable liquidity horizons
2. Constant liquidity horizon and variable position exit costs

The current Committee proposals fall into the first category, which essentially assumes that positions can be managed down over defined periods of up to one year depending on the market liquidity of underlying risk factors. This would correspond to a 'going concern' Pillar 1 capital treatment.

The second approach reflects a 'gone concern' Pillar 1 capital treatment, and aims to model the (stressed) cost of exiting all positions over a single, relatively short, risk horizon. This would closely correspond to the cost of selling a firm's trading portfolio to another market participant in a stressed environment.

The two approaches can be conceptually linked by noting that the premium required by a buyer to take on another firm's trading positions should be at least equal to the subsequent cost of running those positions down over a more extended period afforded by not being a forced seller.

We can see merit in both approaches, at least at a broad level, and would be supportive in principle of a framework which allowed either to be followed. In the forthcoming subsections, we provide more detailed feedback on each.

3.1 Variable liquidity horizons

One of the key challenges for this approach is how market risk factor changes over varying liquidity periods should be applied to the current portfolio. Whilst we would agree that the 'constant level of risk' assumption has led to some non-intuitive results, we do think that a framework for rolling the portfolio forward in time is valuable. For example in the context of the Incremental Risk Charge ("IRC") liquid positions can be defaulted multiple times over the capital horizon. A number of the more recent examples of trading book losses arose not necessarily due to the immediate impact of significant market movements, but rather because of the subsequent difficulty in hedging the resulting positions in what became a stressed market, often with many of the main participants positioned the same way. Our understanding is that this issue also informed the requirement to model hedge slippage costs in the recently introduced Comprehensive Risk Measure ("CRM") capital charge.

This would suggest to us that the alternative proposals where changes are applied instantaneously potentially ignore an important source of risk. We do recognise of course that explicit modelling of the portfolio forward in time is challenging for a variety of reasons, but ruling this approach out at this stage does not seem the right solution.

This would then require guidance on the appropriate rollover assumptions, both in general and for positions falling due during the liquidity horizon. In this regard, we would support the proposed liquidation approach whereby positions are assumed to be fully hedged (or closed) at the end of the liquidity horizon (and not subsequently rolled over) for capital purposes. We would also see value in being able to consider alternative rollover strategies for internal risk management purposes.

Of the other two options proposed for applying risk factor shocks over longer and varying horizons, we would not support the scaling of one-day risk factor shocks. This is highly likely to produce unrealistic, and in many cases arbitrageable, market scenarios, particularly when applied to daily changes derived

from stressed market conditions. The use of a weighted-average liquidity horizon would be operationally problematic to implement, given the possibility that this could change frequently due to portfolio changes.

We would argue that a better simplified alternative to a model where positions are simulated forward in time, and risks are re-hedged (or positions exited) using risk-specific liquidity horizons, would be one where the ES for individual risk types (i.e. FX, interest rates, equity, etc.) are calculated over a common (and fixed) liquidity horizon, for example 10 days as is the case now, and the results are scaled up to account for liquidity requirements in excess of this period. This is similar to the proposal by Morgan Stanley in November 2010, 'A liquidity-sensitive VaR-based framework' that was presented to regulators.

The key benefits from this approach are its simplicity and its similarity to the models already in common use within the industry. However, it has the potential to underestimate risk as argued above, and essentially eliminates cross-asset class hedging benefits as we note in our response to Q6, and therefore ought to be viewed in the context of the fuller simulation approach. We would consider that a framework whereby both approaches were admissible, but where the use of the simplified model would either need to be accompanied by analysis demonstrating that risks were not underestimated by comparison to the more sophisticated model or perhaps be subject to further regulatory scaling, would be a practical proposition. This would then both provide an incentive for firms to further improve this aspect of their risk modelling capability, as well as permitting more simplified approaches where this cannot be easily achieved. In conceptual terms we would see this as similar to the approach adopted by the Committee historically for 10-day VaR, whereby the use of scaled 1-day VaR is permissible if it can be demonstrated to be appropriate relative to the calculation of VaR using 10-day shocks.

3.2 Variable exit costs

Our understanding of emerging industry proposals on this alternative approach are that it is intended to capture the risk of the trading portfolio over a common, relatively short, horizon (e.g. 10 days), as well as the incremental exit costs for any positions which cannot be sold at the bid-offer spread reflected in current P&L reserves. The modelled bid-offer spreads would need to be representative of what one could reasonably expect in a stressed market for the given size of positions (the consideration of stressed market conditions differentiates this from the existing bid-offer reserve computed by Finance).

The use of a common, and short, horizon greatly simplifies the ES calculations. If this were fixed at 10 days, which under the assumptions discussed within the Trading Book Group can be demonstrated to equate to approximately a 30-day (or 6-week) resolution period for winding up a firm, then this would of course be consistent with the current regulatory VaR calculations. This would therefore be more straightforward for firms to implement.

Whilst we can see that there is a certain conceptual simplicity in this approach, and it does avoid some of the modelling challenges associated with differential liquidity horizons, we think that more time could be required to more clearly define how the incremental position exit costs are to be calculated in practice. In particular, we would note that there is significant complexity, and disparity, in the way firms currently calculate bid-offer reserves, and adapting this calculation to stressed markets adds a potentially significant further layer of complexity.

3.3 Conclusion

We fully support the differential modelling of liquidity in the market risk capital framework. We can see merit in both a variable liquidity horizon as well as a variable exit costs conceptual framework, and feel that either could be used to support the aims of the Committee. At this stage, our preference would be for the variable liquidity horizon approach, primarily because we believe this is likely to be more closely aligned with the models which will be used for risk management.

4. Expected Shortfall

Q8. What are the likely operational constraints with moving from VaR to ES, including any challenges in delivering robust backtesting, and how might these best be overcome?

We support the proposal to move from VaR to ES, and recognise that, with suitable calibration of the threshold, this potentially provides a mechanism to achieve a more unified trading book risk calculation framework. In particular, we would hope that risks which previously could not be properly assessed and capitalised through VaR, for example gap risks, could now be integrated into ES models, rather than capitalised through a range of different (and standalone) risk models. An important consequence of this is that it would significantly reduce instances where risks are double-counted, as is the case currently due to the use of different and overlapping risk measures.

Another benefit with ES is that the marginal sensitivities to changes in positions would be more stable than when considering a fixed percentile, at least for the models currently used by some firms. When coupled with the sub-additivity of ES, this would potentially facilitate better risk allocation and capital management at the trading desk level.

4.1 ES calculation

In order to achieve accurate estimates of tail risk, we would expect that a significant number of above-threshold loss scenarios would need to be simulated and averaged over. This is likely to prove problematic for banks which currently use a historic simulation VaR methodology, or some close variant thereof, unless the threshold is set on or below the 95th percentile. In this case, when using 500 days or more of recent historical data, at least 25 loss scenarios would contribute to the ES calculation. We think this will be an important consideration irrespective of whether parametric or non-parametric methods are used to estimate ES.

The dependency of ES on all tail scenarios will naturally lead to additional challenges relating to data cleansing, when viewed against the current VaR requirements. For example, incorrect historical market data for even a single day can potentially have a significant potential impact on ES, where by comparison the 99th percentile VaR estimate may be relatively insensitive to this. To a large extent, firms have already implemented comprehensive data monitoring and cleansing processes, so we would not regard this as a major problem.

However, notwithstanding the above points, we would caution that for firms for which the modelled loss distribution has a very long tail, perhaps due to specific tail risk activities, producing reasonable estimates of ES will prove problematic regardless of the choice of threshold. In this case, fundamental changes in either the risk model itself and/or its implementation are likely to be required. For example, it may then be necessary to move away from a pure historical sampling paradigm towards a Monte Carlo simulation approach, in order to increase the number of tail scenarios that can be generated. As the majority of firms currently use a variant of historic simulation VaR, we think this potential outcome should be taken into consideration given the very significant operational impact (particularly if full revaluation is required). We therefore recommend that the planned QIS specifically investigates the seriousness of this issue to enable an industry-wide solution.

4.2 Backtesting

To the extent that the Committee intends for ES forecasts to be backtested explicitly, we anticipate that this could require significant changes to the operational processes developed for VaR backtesting, which to a large extent are currently focussed around the use of exception counts. Standard techniques for backtesting ES do exist, and would enable the Committee's objective of testing not only the frequency but size of exceptions to be realised. We would note that the power of such testing would be greatly improved by the use of a lower ES threshold (e.g. the 95th percentile). In fact, we would be concerned that the use of a higher percentile would undermine the power of the corresponding backtesting to such an extent that this could lead to repeated outcomes either where poorly performing models were not rejected at all, or

alternatively where reasonable models were rejected too quickly on the basis of statistically inconclusive results.

From a more technical perspective, we would note that the likelihood of observing realised ES outcomes under the assumption that the forecast model is correct (i.e. the null hypothesis) is not, in general, analytically known (e.g. it will be influenced by the kurtosis of the distribution). Therefore, it would instead need to be established empirically, most likely through Monte Carlo simulation. This is quite different from the backtesting of percentile measures, where the distribution of exceptions is expected to be binomial under the null hypothesis. This adds a significant element of additional complexity to the backtesting process. We believe that, as is the case with counterparty credit risk model backtesting, this is warranted and can add significant value to the model testing framework. However, we would caution that it will take time to adapt existing backtesting implementations to deliver this functionality.

As is the case for the existing regulatory backtesting requirement, we assume that the intention would be to backtest the model over a 1-day horizon, rather than the liquidity-adjusted capital horizon.

4.3 Full repricing

We believe that the proposal to require the use of full repricing methodologies for ES is likely to have very significant repercussions from an implementation and cost perspective for some banks, at least in its current form. VaR calculation systems that use a mixture of risk sensitivities and partial revaluation scenarios (e.g. pricing grids) will require significant (re)development to achieve this objective. The associated investment costs and elapsed time required to effect this change at an enterprise level should not be underestimated. This is likely to have some bearing on the planned QIS activities in 2013, where without sufficient lead-time it would be impossible to have full repricing methodologies available.

In fact, we would question whether it is necessary to strictly require the use of full repricing methods. An alternative proposal that could largely achieve the same objectives, and would be more straightforward and less costly to implement, would be to require periodic ES comparisons with a full repricing methodology. Wherever either the results from this comparison, or from the related risk P&L attribution analysis proposed by the Committee, are not of a sufficient quality, then corrective action could be required (or model approval withdrawn for the corresponding business/desk). This would provide firms with more flexibility in their implementation, and in particular enable the benefits from fast, and accurate, approximation techniques to be retained where their effectiveness can be robustly demonstrated. The development of a full repricing environment for model testing would still be very challenging, but at least the deployment costs would then be moderated.

4.4 Impact analysis

We have carried out some preliminary analysis comparing ES calculated using a 95th percentile threshold with 99% VaR for RBS positions, using output from the current regulatory VaR models. As one might expect, the relationship between these measures can vary over time, and will of course depend on the particular portfolio. However, we find that, over the period from Jan 2011 to Jun 2012 for RBS plc regulatory positions, the average ratio between ES and VaR is 0.84 with a standard deviation of 0.08. For Stressed VaR, a similar analysis produces corresponding results of 0.85 and 0.05. Given the stated intention for ES models to be calibrated to a period of stress, the latter comparison is probably more relevant.

In the context of the question as to whether credit default risk should be integrated into the same model as other trading book market risks (refer to section 6), it is perhaps worth noting that this analysis can look significantly different for the current Incremental Risk Charge (IRC). In large part due to relatively concentrated exposures to highly-rated sovereigns, the tail of the loss distribution can be very long and "lumpy". This can lead to ES results which are relatively much higher when compared to VaR results at a fixed percentile. We assume this will be analysed further as part of the forthcoming QIS exercises.

4.5 Conclusion

While we support the move to ES, there are a number of detailed points that would need further discussions prior to a QIS exercise and determination of final rules.

Q5. What are commentators' 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?

For practical reasons, we would support the use of an "indirect" approach to calibrate models to a suitable stress period. Our experience with the estimation of the 1-year historical period for Stressed VaR, taken together with some of the additional ES requirements proposed by the Committee (and in particular the use of full revaluation), leads us to the conclusion that the "direct" approach may be unworkable. However, we do not agree with the specific "indirect" method proposed by the Committee as it stands. In particular, we note that the proposal does not deliver a stressed ES measure, but rather a maximum stressed loss scaled by the ratio of two ES measures.

We would prefer a variation of the "indirect" method whereby the unstressed, or current, ES is scaled by the ratio of stressed versus unstressed ES calculated for a reduced set of core risk factors. In this respect, we endorse the view presented in the industry feedback consolidated through ISDA, and the corresponding proposed stressed measure ES_S defined by:

$$ES_S = ES_{FC} \times \frac{ES_{RS}}{ES_{RC}}$$

where:

ES_{FC} is the ES based on the full set of risk factors in the current period,
 ES_{RC} is the ES based on a reduced set of risk factors in the current period, and
 ES_{RS} is the ES based on a reduced set of risk factors in a stressed period.

To some extent, we can think of this as applying a variable adjustment factor to a 'point-in-time' outcome to convert to a 'downturn' (or stressed) outcome.

In our view, this would have the following benefits:

- It is closely tied to the ES calibrated to current market conditions, which is likely to be more extensively used for internal risk management;
- The requirements for backtesting, and more generally model validation, of the full risk measure are much clearer, as the expected statistical properties of the model forecasts are well defined. For example, VaR exceptions at the 95th percentile are expected to occur on average 5% of the time;
- The calculation and identification of an appropriate stressed ES is much more tractable with a reduced set of risk factors, from both a computational perspective as well as with regards to compiling a complete set of market data histories.

Should this approach be followed in conjunction with setting the ES threshold at 95%, we believe that the 'use' test could then largely be retained (perhaps subject to the outcome of discussions on liquidity and diversification). Although we understand that regulators are considering de-emphasising this requirement, we do think that maintaining a close link between the models used for internal risk management of the firm and those used to calculate capital is very important to ensure that banks do not become subject to perverse risk management incentives.

5. Diversification

Q6. What are commentators' views on the merits of the desk-based and risk-factor-based aggregation mechanisms to deliver the Committee's objectives of constraining diversification benefits?

The requirement for diversification and hedging benefits to be reliable in a stressed market environment is reasonable, and we do acknowledge that correlations at the broad risk factor level have the potential to develop differently in future market crises to those we have seen in the past. For example, the negative correlation experienced between major currency interest rates and credit spreads as a result of coordinated central bank intervention to cut interest rates during the most recent crisis may not necessarily be expected to transpire in all future episodes of market stress.

However, the approach presented by the Committee in which asset class correlations are prescribed by the regulators does seem to present a number of difficulties. For instance, it is not at all clear how one would determine whether a given portfolio is 'long' or 'short'. Particularly for portfolios which are not directional, but also more generally for complex portfolios containing options and other structured products, we do not see an obvious way in which this could be done.

One possible solution would be to ignore any directional effects. However, this would mean that, whilst diversification benefits would still be possible, direct cross-asset class hedging would not be recognised. This would potentially be problematic, particularly where the determination of asset-class for a particular risk was not clear. Given current market circumstances, a very relevant example would be local currency sovereign debt, which could conceivably be classified as either interest rate risk or credit risk. Under the former classification, hedges using interest rate swaps would be recognised within the model, whereas in the latter case they would not. In practice, these assets may behave more like interest rate risk at certain times, but like credit risk at other times — the credit rating of the sovereign can be an important determinant of this.

If firms were to be given some flexibility as to how risks were assigned to the designated risk types (or asset classes), subject to the appropriate qualitative standards, then this may largely mitigate concerns of the above type. However, it would still offer very limited recognition to certain types of cross-asset class downside risk hedging activities. For example, the hedging of long equity exposure via the purchase of credit protection on the same name, which although not perfectly correlated should offer a high degree of risk mitigation in extreme scenarios.

Lastly, it is not clear to us how regulators would be able to consistently calibrate industry-wide correlation parameters, and that this would not lead to unintended consequences. For example, by giving insufficient capital benefits to more diversified portfolios relative to highly concentrated portfolios, or (to the extent that regulatory capital is a constraint on business) strongly influencing firm-level decisions on the optimal portfolio mix across different sub-portfolios/risk types.

Given these concerns, we are uncomfortable with this general approach. We are aware of an alternative proposal in which the level of diversification benefit would be set on a sliding scale, from full diversification as captured by the internal model at one extreme to no diversification at the other.

$$Capital = \sum_{i=1}^N IMCC(C_i) \times (1 - \alpha) + IMCC(C) \times \alpha$$

The corresponding scalar α , between 0 and 1, would be determined by regulators based on their assessment of the extent to which asset class correlations were reliably modelled within the bank's ES measure. We think this idea may be more practicable as a way of limiting diversification benefits, although it would need to be developed further, particularly so that the standards required to use internally modelled correlations (i.e. achieve $\alpha = 1$) were clearly articulated.

Conclusion

Whilst we understand concerns with the recognition of diversification benefits which may not be fully robust to a stressed environment, we do not support the Committee's proposals for correlations to be specified by regulators. Instead, we would contend that firms should continue to be given flexibility to determine correlations for use within their internal models, but that these should be subject to strengthened validations standards in order to address the Committee's concerns. In particular, firms would need to be able to demonstrate that correlations are appropriately calibrated to stress conditions.

6. Credit risk

Q7. How can regulators ensure robust supervision of integrated market and credit risk modelling? In particular, how would an integrated modelling approach affect other elements of the proposed framework (e.g. the choice of the quantile parameter for ES, the P&L attribution and backtesting processes, etc.)?

One of the main challenges with modelling defaults in particular is that, because occurrences are relatively rare, it requires many simulation paths to produce reasonable estimates for the loss distribution. Therefore, any joint model of market and credit risk will need to overcome the problem of repricing the full trading book, and all associated pricing risks, a very high number of times. Given the intention to require firms to use full repricing methodologies, this is likely to prove very expensive and time-consuming operationally. The usability of the models for internal risk management purposes, where intraday calculations and near-time availability of results can be important, may be compromised.

Furthermore, the joint modelling of market and credit risk factors involves significant complexity, particularly if the intention is, as we believe it should be, to simulate realistic non-arbitrageable future market conditions. We think it would be a mistake to require a joint model framework if it subsequently leads to overly simplistic modelling of market risk factors.

For reasons strongly linked to these, RBS chose to implement separate price and default risk measures as part of its CRM model developed during 2010-11.

This could suggest a need to retain separate capital charges for market risk and (traded) credit default risk. However, we do recognise that there are benefits to be obtained from a suitably specified joint modelling framework, particularly in helping to develop additional understanding of the interplay between market and credit risks. From an internal risk management perspective, we see great value in being able to consistently view market and credit risks across the entire trading book. Therefore, we would be inclined to suggest that the framework be left open at this stage to allow either separate or joint modelling of these risks. Clearly we would expect supervisors to impose additional requirements before accepting the use of joint models, not least as one could expect these to facilitate some degree of capital saving through additional diversification. This would then provide reasonable incentives for firms to continue to develop their modelling capability in this area, without requiring significant subsequent change in the modelling rules to recognise such improvements.

To some extent this mirrors the treatment of CRM models, where either integrated or separate modelling approaches can be followed, although importantly the overall standards (i.e. 99.9% loss over 1-year horizon) are identical. In this latter respect, we have argued in previous responses that ES models should be calibrated using a 95th percentile threshold. Taken together with the Fundamental review proposals to use risk-specific liquidity horizons (with a minimum of 10-days) and an assumption that risks are fully liquidated at the end of this period, this could imply a much lower level of capital to be held against default risk. It is currently unclear to us what approach is intended to be taken with regard to scaling of risk measures to produce the overall capital. However, as a general point we do think that the modelling of joint default outcomes for a portfolio of credit positions with liquidity horizons up to 1-year is more tractable when considering ES at the 95th percentile than VaR at the 99.9th percentile. The former measure still provides an estimate of the average '1-in-20 year' loss, at least for the most illiquid part of the portfolio. The use of a stressed calibration will further increase the severity of the modelled losses.

In the event that separate models are permitted, we would strongly advocate that potential losses from credit rating migrations be incorporated into the price risk measure and not the default measure. This would avoid much of the double-counting that exists in the current regulatory model framework between (Credit) VaR and IRC models.

We do not consider that a joint model would impose significant extra difficulties in terms of either P&L attribution or backtesting.

Conclusion

We would like to see a framework which would allow for, but not require, the joint modelling of credit risk and market risk. In order to achieve this, the modelling standards need to be aligned, and in particular the threshold at which ES is to be calculated must be consistently specified. This would then permit separate models to be used by firms who have yet to develop this capability, and provide a more seamless transition from the current regulatory modelling rules in which separate VaR and IRC models are used.

7. Model oversight and approval

Q4. What are commentators' 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 positions? Are there alternative classifications that might deliver the same objective?

We are supportive of a more granular model oversight and approval process, and the related requirements to demonstrate both adequate backtesting and P&L attribution analysis at a desk level on an ongoing basis. In the case of backtesting, we would note that this analysis is already performed on a regular basis, and therefore should not present any additional difficulties.

The proposed portfolio level approach seems reasonable, under the assumption that the size and number of portfolios remains broadly as exemplified by the stylised example given in Table 4 of the Fundamental Review. We would see potential dangers in extending this approach too far down the organisational hierarchy, both in terms of the increased operational and administrative overheads as well as the greater 'noise' that might be expected in the results.

7.1 Mitigating 'cliff' effects

In common with many other commentators, we are concerned about the potential 'cliff' effect arising from switching from internal model to standard rules for a particular desk. We would prefer to see a more gradual transition between the two, and feel that this could best be achieved by a variation on the proposed approach. A measure of model backtesting performance over the most recent period could be used to define a variable scalar between 0 and 1, with a higher value representing poorer performance. Capital would then be determined as a weighted combination of the internal model and standard rules using this value.

This is similar in many ways to the existing use of the number of VaR exceptions over the previous 250 business days to determine the need for, and size of, an additional capital buffer. In fact, we can easily generalise the currently specified criteria to allow more direct application to ES backtesting. For example, in the table below we have simply transformed the number of exceptions into an approximately equivalent p-value range, where the p-value represents the probability of the observed model backtesting outcome under the null hypothesis (i.e. assuming the model is correct). The table is only for illustrative purposes.

P-value from	P-value to	Scalar ω
0.1	1.0	0.0
0.05	0.1	0.4
0.01	0.05	0.5
0.005	0.01	0.65
0.001	0.005	0.75
0.0005	0.001	0.85
0.0	0.0005	1.0

If this approach were to be adopted, we would of course expect further refinement of both the p-value ranges as well as the corresponding scalars as part of the consultation process.

For the purpose of overall consolidation, it would be necessary to scale the corresponding desk risk:

- To produce $1 - \omega$ times marginal ES for aggregation of ES with all other (eligible) desks; and
- To produce ω times marginal standardized approach for aggregation of standard rules based capital.

Furthermore, with this approach in place, there would clearly be no need to retain the current incremental capital multipliers introduced in the event of progressively worse model backtesting performance.

7.2 P&L attribution

Whilst we agree that P&L attribution analysis can be an important tool to build confidence in the degree to which the internal model is able to explain the actual P&L, we do see some areas where further clarity would be required. This mainly relates to the treatment of idiosyncratic risk factors which, depending on the methodology, may not be modelled directly but instead captured via, for example, a multi-factor regression. The main purpose of such a model is then to forecast the potential variability of these factors over the appropriate horizon — the model may legitimately have limited ability to forecast the future level of these factors, particularly where they are highly idiosyncratic (i.e. in a regression sense, the R^2 is low).

More generally, we assume that any source of daily P&L which is either explicitly dropped from the model, for example a risk factor deemed to be unmodellable, or which could not reasonably be expected to be captured by the overall model framework, i.e. P&L due to intraday trading, fees, commissions, etc., would be excluded from the P&L attribution analysis.

These types of complications strongly suggest to us that an advanced P&L attribution framework would be required, giving the ability to stage a controlled P&L 'roll' by incrementally adjusting each of the components of risk. For example, for an FX derivatives portfolio, first changing FX spot rates, then FX implied volatilities, then interest rates, etc. In certain instances, it is possible that this sequencing of risk factor changes could create intermediate market states which are inconsistent and arbitrageable. Notwithstanding these and other related technical difficulties, we would note that, for banks that do not currently have this full capability, this is likely to be a very significant undertaking and therefore will take time to implement.

7.3 Conclusion

We agree with the proposals to extend model approval to desk level, and the broad outline of how model performance would be assessed. In terms of details, the definition of a “desk” remains a key question that would need to be clarified. We would also welcome further detailed discussion on how the P&L attribution analysis would work.

8. Standardized approach

Q3. What are commentators' views on the proposed regime to strengthen the relationship between the standardised and internal models-based approaches?

We fully support the intention to bring these approaches closer together, and would agree that one of the main impediments to their wider use has been the failure of standardised methods to even approximately capture the risk of hedged derivatives portfolios. For example, we think that many of the concerns around the use of a floor to the recent CRM capital charge, based on a percentage of the corresponding standardised method, would have been mitigated if this had more adequately captured hedging benefits, and in particular rewarded rather than penalised (effective) hedges.

We agree with the requirement for all banks to calculate results for the standardised approach, irrespective of whether they have approved internal models. This should provide a further means of challenging internal models where rank ordering of results from standardised and internal models-based approaches across firms shows up inconsistencies'. It should also enable ongoing evidence to be gathered to help ensure that the alternative methods can continue to be closely calibrated, and lead to the early identification of portfolios for which the standardised approach is clearly inadequate. We would note that any such comparison of models will need to be carefully managed to ensure that the conclusions drawn are reasonable, and so that firms are given ample opportunity to explain any apparent anomalies in their internal model results.

From an operational perspective, we can appreciate this will also facilitate plans to manage the approval of models at a desk rather than legal entity level, allowing a more seamless switch between advanced and standardised approaches in the event of poor model performance (and the converse where the required model improvements have been demonstrated).

Our support for the use of capital floors on internal models based on a percentage of the standardised approach would be highly dependent on both the choice of the floor, and the implementation of a significantly improved, and more risk sensitive, standardised approach. Otherwise, we would be very concerned that the floor would become the dominant measure, and thereby potentially provide perverse risk management incentives where capital is not aligned with economic risks. This has been the experience of many firms in relation to the CRM capital charge. We also think that this could act as a disincentive to the continued development and improvement of risk models.

Q9. Which of the two approaches better meets the Committee's objectives for a revised standardised approach?

Of the two approaches described in the Fundamental review (and in the case of the fuller risk-factor approach, the more detailed associated document entitled 'Fuller risk-factor approach: A proposal for its further specification'), we would tend to prefer the fuller-risk factor approach. Our initial impression from studying the two models is that this approach is likely to be more risk sensitive, and better able to quantify the risks from non-linear positions (e.g. options). We would regard this as similar to a variance-covariance risk model, with correlations defined by a restricted hierarchical framework, and some adaptation for non-linear risks. Therefore, whilst the calibration of model parameters in particular has yet to be fully defined, we can see some promise in this approach.

We would note that the need for a firm to use its own pricing models to implement this model, at least for some product types, would appear to contravene one of the Committee's stated objectives. However, we believe this is warranted in order to produce an appropriately risk sensitive charge, in particular in the case of non-linear instruments where the notional equivalent value (or position size) at current market rates can be a misleading indicator of risk (at least when used in isolation).

The level of sophistication of this approach is clearly higher, and whilst this should help to close the gap to a large extent with results produced using firms own internal models, we do appreciate that this may create operational challenges for some of the smaller firms to which these rules are also intended to apply. In this regard, we would perhaps suggest that a feasible option would be to consider retaining the current standardised rules as a further fallback. In this way, firms that were unable, or unwilling, to implement the fuller risk-factor approach would have an alternative, albeit one that is less risk sensitive and likely to generate higher capital requirements. However, firms using internal models would then be required to implement only the fuller risk-factor approach, both as a fallback in case of model approval being withdrawn, as well as a means to more effectively benchmark results across the industry.

More generally, we would suggest that further consultation on this topic would be welcome, both to improve industry understanding of the finer details of what is being proposed, as well as to allow the results of further analysis to be incorporated into the development and calibration process. Realistically, we would only expect to be able to give more definitive feedback once a more detailed model description and an initial calibration proposal have been provided.

Conclusion

On balance, and taking into consideration the intentions of the Committee with regard to the wider use of standardised model charges for both industry benchmarking as well as a potential capital floor to internal models, we prefer the fuller standardised approach.