185 Spadina Avenue
Toronto, Ontario M5T 2C6 Canada
Tel 416 217 1500
Fax 416 971 6100

Basel Committee on Banking Supervision
Bank for International Settlements
Centralbahnplatz 2
CH-4002 Basel
Switzerland

August 24, 2012

Dear Members of the Basel Committee on Banking Supervision:

Algorithmics, an IBM Company ("Algorithmics") appreciates this opportunity to comment on the "Fundamental review for the trading book" consultative document ("the Review") issued in May 2012 by the Basel Committee on Banking Supervision ("the Committee"). As a provider of risk-related consulting, software and services to more than 350 financial institutions worldwide, we have undertaken discussions of the Review with a wide range of banks that will be directly impacted by the proposals. For your reference, we provide a corporate overview of Algorithmics in Annex F.

Overall, we agree with the Committee's key observation that the current market risk framework lacks coherence as it does not offer a single overarching view of how trading risks should be categorized and capitalized. We also referred to this issue of a "Mosaic of Measures" in our response to the "Revisions to the Basel II market risk framework" in 2009[1]. At the time, we believed such a patchwork of measures would not only fail the use-case but also lead to real consequences in terms of operational challenges for banks as well as supervisory and validation challenges by regulators. We applaud the Committee on adopting similar views and undertaking a fundamental review to address these shortcomings.

We also agree with the Committee that an integrated framework does not necessarily imply a single, unified model for all risks (i.e., a one-size fits all approach) and we see nothing wrong, at a high level, with Steps 1 and 2 in Figure 1 of the Review. In fact, having the ability to withdraw internal model approval at the desk level as opposed to the all-or-nothing current approach is a sound way, in our opinion, to allow banks to continue to use robust models where they work while providing them with an incentive to improve their risk management approaches in areas where they may currently fall short.

However, we disagree with many of the prescriptive modelling requirements in Step 3. Requiring banks to separate out risks and recombine them later through top-down aggregation functions, adding standardized surcharges for jumps in liquidity premiums, calculating add-ons for risk factors deemed not modellable, and adding separate capital charges for default and migration, only serve to replace the current patchwork of measures framework with another.

[1] http://www.bis.org/publ/bcbs14849/ca/algorithmics.pdf

Moreover, a detailed study of the modelling suggestions in Step 3 has highlighted the presence of many issues, concerns and deficiencies. These issues are discussed at length in our response to the Committee's questions below as well as in the Annexes attached to this letter. **The key finding is that many of the suggestions in the Review will create a framework that is rife with the possibility of regulatory capital arbitrage** (i.e., the ability to manipulate regulatory capital requirements without changing the true economic risk profile). These arbitrage opportunities arise because the Committee, despite its best intentions, has not actually delivered an integrated framework for assessing risks.

In our opinion, designing an integrated framework should mean that if a risk is deemed to be internally modellable (i.e., Step 3), then it should be allowed to be modelled in a manner consistent with other risks that are also deemed to be internally modellable and that are inextricably linked. Only in this manner can you avoid the manipulation of any single risk measure.

To this point, perhaps a *fundamental* review for the trading book was not entirely necessary. The Committee has already sanctioned the use of a truly integrated framework for the internal modelling of Counterparty Credit Risk ("CCR") and more particularly, for measuring Effective Expected Positive Exposure ("EEPE"). Applying many of those modelling elements towards market risk measurement is, in fact, the ideal integrated framework the Committee should be advocating. In what follows, we describe in detail how the CCR internal models framework can be merged with a Market Risk ("MR") framework to achieve a single Unifying Framework. We explain clearly how such a modelling framework solves many of the problems present in the Committee's current proposal.

A Unifying Framework

MR and CCR are inextricably linked; losses in portfolio values lead to MR while gains in portfolio values feed CCR. Thus, a consistent approach to modelling portfolio value distributions through time will provide the foundation for proper estimation of both MR and CCR. Historically, the synchronization of both risk measures was deemed irrelevant given the relative differences in time horizons.[2] However, the Committee's current proposal to adopt liquidity horizons of up to 1y in length for MR has now introduced the opportunity for consistent overall simulation of both MR and CCR. The table and sections below describe in detail what such a Unifying Framework could look like and how such a framework can solve many of the problems present in the Committee's current proposal.

[2] MR traditionally took a 10d horizon while the Effective EPE measure for CCR took a 1y horizon.

	Suggested Market Risk (MR) Internal Model	Internal Model for CCR (EEPE)[3]
1) Market Risk Factor Calibration	- Calibration to occur over stressed period (spanning 1 year of historical data).	- Calibration to occur over current period as well as stressed period (spanning 3 years of historical data in either case).
2) Scenarios & Time Steps:		
Risk Factors	- Scenarios based on all market factors impacting instrument valuation (i.e. FX, IR, etc.)	- Scenarios based on all market factors impacting instrument valuation (i.e. FX, IR, etc.)
Scenario Type	- Multi-step scenarios (distribution of risk factor values at multiple time steps)	- Multi-step scenarios (distribution of risk factor values at multiple time steps)
Measure	- Historical measure	- Risk Neutral or Historical measure
Time Horizon	- Time Horizon of up to 1 year	- Time Horizon of 1 year for EEPE and full maturity profile for Effective Maturity
Time Steps	- No restrictions on placement of time steps but should include (at minimum) the required liquidity horizons: 2w, 1m, 3m, 6m and 1y	- No restrictions on placement of time steps
3) Simulation & Valuation	- Full forward valuation of portfolio	- Full forward valuation of portfolio
4) Value Transformation	- Values converted to Profits and Losses ("P&L") by comparing Value under each scenario and time step to current value.	- Values converted to Exposure taking into account Netting relationships
5) Mechanical Adjustments	- Settlement (expired instruments) - Liquidation (at designated liquidity horizon) for instruments not expired	- Collateral (Margin Risk) - Ratchet and Time-weighting of profile
6) Measure	- Expected Shortfall (at 99%, 1y)	- Expectation for EEPE measure - VaR on Full EE profile for CVA
7) Backtesting	- Attribution - Backtesting	- Full distribution of Exposure - Full distribution of Risk Factors

Table 1: Unifying Framework for MR and CCR

[3] Material here is from Annex 4 of the document "Basel II: International Convergence of Capital Measurement and Capital Standards: A Revised Framework - A Comprehensive Version" (available at www.bis.org/publ/bcbs128.htm) as well as from Section II.A of the document "Basel III: A global regulatory framework for more resilient banks and banking systems" (available at www.bis.org/publ/bcbs189.htm).

Steps 1-3 represent the key modelling steps (i.e. Calibration, Scenario generation and Simulation/valuation) in the Unifying Framework. As shown in Table 1, these steps are identical for both MR and CCR measures, which leads to many benefits:

- This consistency will promote real advantages to banks in terms of operational deployment, input data consistency, data mapping, model development and analysis.
- From a regulatory standpoint, validation will need to occur only once for both Risk measures, thereby streamlining and strengthening the supervisory validation process.
- Allowing the same framework for both measures will provide even greater incentive for banks to improve their risk methodologies as there is greater re-use of approaches.
- Enforcing a consistent, unifying framework between measures reduces the potential for regulatory arbitrage as banks would not be able to develop entirely different approaches for each risk measure.

Steps 4-6 represent slight differences regarding "mechanical" manipulations of the value profiles of the portfolios. Since there is very little modelling choice here, these steps are not the subject of ongoing, intense regulatory supervision.

Step 7 represents the "proof" point that the models work. In both MR and CCR the Committee has expressed the need to check both the pricing models as well as the risk factor distributions. However, the need is expressed slightly differently in each case. The pricing models are tested under Attribution in MR and by looking at the full distribution of Exposures in CCR while the risk factor distributions are encapsulated in Backtesting for MR and by looking at the full distribution of Risk Factors in CCR. There are clear opportunities to synchronize backtesting efforts here for both MR and CCR. In the meantime, our responses to Questions 7 and 8 focus on how to incorporate robust and sensible backtesting procedures for MR.

In what follows, we describe each section in more detail and, where appropriate, reference other sections in our response or Annexes that highlight how the proposed Unifying Framework addresses many of the shortcomings in the current Review, including the elimination of regulatory capital arbitrage opportunities.

1) Market Risk Factor Calibration

While both MR and CCR approaches emphasize the need to include a stress period, the period selected may be different between the two measures as MR focuses on the left tail of the distribution (losses) while CCR focuses on periods of stressful movements in credit spreads for the portfolio. In reality, periods of stress for credit spreads will most likely coincide with stressful movements in other risk factors meaning that both MR and CCR measures are likely to have common stress periods. Regardless, the benefit from a harmonized approach is not derived solely from the fact that they may share the exact same stress period but rather, from the fact that they share the same risk factor representation, historical data, calibration routines, risk factor simulation models and mapping of risk factors to positions.

1.a) Top-down aggregation and stressed parameters
The internal model for CCR requires banks to consider the Max(Current EEPE, Stressed EEPE). The current proposal for Market risk requires banks to consider Max(Integrated ES, Aggregated ES). Both ES measures are based on stressed parameters. However, Aggregated ES requires banks to break up their risks into general risk factor classes (i.e. IR, EQ, FX etc...) and aggregate them back together using an aggregation function and regulatory specified correlation between the risk factor classes. This is presumably to limit the level of diversification. Our analysis indicates there are many issues with the suggested use of the aggregation function, not the least of which is that it leads to regulatory capital arbitrage. These issues, as well as arbitrage examples, are presented in Annex B and in our response to Question 6 below.

Conceptually, since the Unifying Framework is already calibrated to a stressful period, there should be little need to overlay additional requirements to limit diversification benefits. For this reason we believe the top-down aggregation can be removed with little impact. However, should the Committee desire some level of control over inter-group diversification benefits, a more direct approach would be to move away from the use of arbitrary, top-down aggregations and allow regulators to "adjust" the bank-provided risk factor correlation matrices directly. Annex C describes: i) how to "roll up" an aggregate correlation matrix from a risk-factor level correlation matrix, ii) how to overlay regulatory views of inter-group correlations at these aggregate levels and iii) how the adjusted, aggregate correlations can be cascaded back to the detailed risk factor correlation matrix in order to arrive at a new, adjusted matrix that can be used in a full simulation model (without the need to split up risks and then recombine them artificially).

With this suggested approach in place, capital can be calculated as Max(Integrated ES, Regulatory Stressed ES) where both measures are calculated using the full simulation approach[4]. Integrated ES would use the bank-derived correlation matrix and the Regulatory Stressed ES would use the bank-derived correlation matrix adjusted with a regulatory overlay of stressed parameters. In both cases, the original bank-derived correlation matrix will be based on a stressed financial period. This will address many of the issues related to the top-down aggregation proposal while still allowing regulators to minimize the extent of diversification.

2) Scenarios & Time Steps
The Review discusses allocating liquidity horizons by major risk class. To accommodate varying liquidity horizons (and potential rollovers), the Committee has suggested the use of an "instantaneous shock" ("IMC") approach for simulation wherein the shock sizes reflect the length of each liquidity horizon. Based on our analysis, it is our opinion that the Committee's reasons for suggesting an IMC approach are misguided and will ultimately lead to a significant potential for regulatory capital arbitrage. Annex A describes our reasoning and findings in this area.

To overcome some of the challenges and inconsistencies brought about by the Committee's suggestion, the Unifying Framework makes use of Multi-Step Monte Carlo ("MSMC") scenarios as is currently done in the CCR framework. All risk factors are simulated on all simulation time

[4] As outlined in our response to Question 6 and in Annex C, the Committee may actually want to specify a positive and negative stress on correlation levels to avoid providing benefit to directional portfolios.

steps out to a horizon of one year. Adopting this approach and removing the top-down aggregation restrictions (discussed in 1.a above) results in banks not having to arbitrarily split up risk factors and scenarios. In order to amalgamate varying liquidity horizons properly, the MR measure would only need to be considered at the 1 year horizon. This point is further described in Annex A.

3) Simulation & Valuation

We are of the opinion that "a swap is a swap is a swap". In other words, regardless of the risk being measured (MR, CCR or something else) , the definition of the swap doesn't change and neither should the valuation approach. Generally speaking, practitioners approach MR and CCR from a completely different mindset when it comes to valuation practices. The short time horizon for MR leaves practitioners with the impression that extreme precision and sophistication in valuation routines is warranted and the use of IMC scenarios generally allows for this precision without affecting overall computational performance. In many cases, front office models are used here. The models may focus excessively on arriving at the best price today (or under an IMC approach), at the expense of ignoring many key path dependent elements[5] which would be present in a MSMC simulation. As we show in Annex A, using "precise" models in an IMC approach actually produces a fall sense of precision.

CCR, on the other hand, is often seen as an exercise in simplifying valuation models and routines in order to address the larger computational requirements brought about by MSMC. The fact that time horizons for MR are now being extended to match more closely CCR indicates that there is a sensible argument to be made for the natural synchronization of valuation models for both disciplines. Ideally, a happy medium should be achieved whereby the same models are used for both MR and CCR. Clearly, the models must be precise enough to pass the Attribution tests the Committee envisions while still incorporating many of the path dependent elements which will lend credibility to its valuation profile in a MSMC context. In the end, the focus should not be on blindly ensuring penny-wise accurate pricing to the detriment of ignoring all else, but rather ensuring the models cater to, and properly handle, all the driving factors that can affect the risk profile. From a regulatory capital standpoint, it is better to be generally right than precisely wrong. This consistency in valuation and simulation between MR and CCR will serve to streamline the regulatory review requirements and lead to tangible operational benefits on the part of financial institutions.

4) Value Transformation

In both MR and CCR cases, the simulated values need to be transformed into a relevant risk measure. For CCR, the values are converted to exposure figures taking into account netting relationships. For MR, the values are converted into P&L figures by comparing the values under each scenario and time step to the current value. While CCR can use the point-in-time value (i.e. the value of the portfolio at a given time step) for the purposes of calculating exposure, the P&L figures must be derived with "all-in" values (i.e. values that not only reflect the value of the

[5] Some path dependent elements include the modelling of rate resets, knock-outs, early exercise, time decay and pull-to-par effects.

portfolio at a given time step but also include things like cash amounts that the instrument may have spun off during the time interval). These transformations are relatively straightforward to specify and model once the full forward valuation is completed. To this end, the value transformations can be treated as a post-processing of the simulated results. In either case, the transformations are prescriptive in nature and require little in the way of modelling judgement. Regulatory supervision should be minimal at this point in the Unifying Framework.

5) Mechanical Adjustments

In the case of CCR, the exposure profiles are adjusted mechanically to account for collateral as well as the explicit regulatory rules for modelling rollovers on the profile through a ratchet and time weight function[6]. In the case of MR, the P&L figures are adjusted to reflect both settlement and liquidation. These concepts are described in more detail in Annex A but, at a high level, they work as follows. Settlement applies to instruments that expire before their liquidity horizon. Contrary to the Committee's statement in Section 4.5.5, these instruments will, in fact, have capital assigned to them via this settlement adjustment. Thus, the Committee's assertion that IMC scenarios are required is incorrect. Liquidation applies to all other instruments that do not expire before their liquidity horizons. The liquidation procedure also offers a direct and intuitive approach to modelling the capital add-on for jumps in illiquidity that the Committee refers to in Section 3.3.2. Annex A, Example 3 shows how the illiquidity add-on can be modelled. Once again the mechanical adjustments here are prescriptive (i.e. very little in the way of modelling judgements is needed) and should not be a major focus for overall regulatory supervision.

6) Measure

CCR and MR are primarily distinguished by the fact that CCR relies on an expectation while MR is a tail-based measure (i.e., VaR or ES). All else being the same, this would imply that MR estimation would require many more scenarios to converge than CCR. However, adopting the Unifying Framework for both CCR and MR would not present additional computational challenges above those of CCR. This is due to the fact that regulatory acceptance of internal CCR models already requires these models to converge to tail-based measures of exposure via the use test requirement[7] . In particular, "*the use test is satisfied if a bank uses other counterparty risk measures, such as peak exposure or potential future exposure (PFE), based on the distribution of exposures generated by the same model to compute EPE*."[8] Given this, there should be little in the way of differences in computational-scale between MR and CCR when adopting the Unifying Framework for both.

It should also be noted that the current discussion of VaR versus ES for MR should not be the focus in our opinion. Since both are simply measures on a final distribution, it is the quality of

[6] In regulatory parlance, the Value profiles are converted to Exposure profiles ("E"). The expectation of the Exposure profiles is then taken ("EE"). The ratchet function is used on top of EE to derive Effective EE. The time-weight function is then applied to Effective EE to calculate Effective EPE.

[7] Please see section titled "Use test" under the Qualifying standards on CCR Management in Annex 4 of the document "Basel II: International Convergence of Capital Measurement and Capital Standards: A Revised Framework - A Comprehensive Version" (available at www.bis.org/publ/bcbs128.htm).

[8] Paragraph 52 in the above reference.

the final distribution that matters, the explicit measure chosen from the distribution is of secondary importance. Our response to Question 8 below provides more detail. Regardless, since the Unifying Framework focuses on deriving the most realistic distribution possible, it is agnostic to the final measure chosen. Both VaR and ES should work just as well in the framework.

Having presented our recommendation for a Unifying Framework for both CCR and MR, we now respond to many of the Committee's questions. The responses will highlight the potential problems arising from various elements in the Review and discuss how the Unifying Framework may be used to address such issues.

Responses to Committee Questions

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

We agree with several elements of the Committee's proposal to capture market liquidity risk through the use of liquidity horizons:

- The general definition of a liquidity horizon is sensible and sound (i.e. the length of time required to sell a financial instrument or hedge all its material risks, in a stressed environment, without materially affecting market prices).
- Less granular, more generic liquidity horizons (5 buckets) limit the amount of complexity while still allowing enough risk sensitivity to properly allocate regulatory capital.
- The use of a regulatory floor for liquidity horizons will limit regulatory capital arbitrage and ensure consistency of measurement and approach between banks.
- Extending the liquidity horizon to account for endogenous effects is perhaps the most reasonable approach.

The key challenge, of course, will be assigning liquidity horizons on the basis of general risk factor classes. While this approach is easy to understand, it suffers from the same limitations as the Committee's proposal to use a top-down aggregation approach (formula 1 in the Review) namely; how to assign risk factor dominance to hybrid instruments. Many common trading instruments face this challenge. For instance, assuming EQ, FX and IR risk factor classes are all assigned different liquidity horizons then:

- Are convertible bonds given an IR or EQ horizon?
- Is a portfolio of equities in different currencies given an EQ or FX horizon?
- Is an FX Swap given an FX or IR horizon?

Not providing clear guidance in this area (perhaps even a definitive set of rules by the Committee) will open the door for regulatory capital arbitrage, wherein the selection of liquidity horizons may be strongly influenced by regulatory capital savings as opposed to an appropriate consideration of the key risks.

In addition to the challenge above, we also disagree with various elements being proposed by the Committee as they relate to the modelling of market liquidity risk. We provide the following recommendations to circumvent the shortcomings:

- Requiring capital add-on for jumps in liquidity premia is reasonable given the experience of the past 5 years. However, the charge should be applied as a standardized capital charge only when that particular instrument/portfolio is being handled in the standardized model. Otherwise, the charge should be handled consistently in the full-simulation approach used for internal models. Under our proposed Unifying Framework, such a charge can be incorporated appropriately as an adjustment applied to the final liquidation value of the instrument at its liquidity horizon. Annex A, Example 3 shows how such a charge can be included. This approach is simple to comprehend, implement and audit. In this case, regulators will be required to provide the supervisory haircuts to be treated as liquidation value Bid-Ask spreads at the liquidity horizon of the instrument.

- The presentation of various options for modelling varying liquidity horizons in the market risk metric (Annex 4, Section 2) is somewhat overreaching and unnecessary. It is overreaching in the sense that, as pointed out in the Unifying Framework section above, the Committee has already accepted a full MSMC approach through a one year horizon for measuring CCR. Prescribing an entirely different, and restrictive, set of modelling rules for MR ignores the Committee's earlier work, ignores the implicit link between MR and CCR and opens the door for regulatory capital arbitrage wherein modelling approaches for MR may be adopted differently from those for CCR simply for the purposes of regulatory capital savings. The discussion of various options is also unnecessary. As we have shown in the Unifying Framework above, using a full MSMC solution (where risk factor distributions are determined at multiple time steps) and allowing the portfolio to be liquidated at its liquidity horizon (with the resultant gains/losses settled into cash) allows a full market risk distribution to be calculated at the end of the 1 year horizon, properly incorporating varying liquidity horizons (without the need for arbitrary simulation rules). This point is further explained in Annex A.

- The Committee considers the use of "instantaneous shock" scenarios to avoid modelling explicit rollover assumptions (section 4.5.5). The need for rollover assumptions is both erroneous and contradictory to the Committee's own definition of Liquidity Horizons. Annex A describes the issue in more detail. The use of instantaneous shock scenarios not only ignores pertinent instrument/portfolio behaviour related to such elements as time decay, path dependency, mean reversion, settlement and knock-outs, it also promotes regulatory capital arbitrage. Example 6, in Annex A provides a simple illustration of regulatory capital arbitrage when using instantaneous shock scenarios. Our recommendation is to allow full forward simulation of the portfolio as is consistent with

the description of the Unifying Framework above. A consistent, forward simulation will capture all elements of portfolio behaviour and prevent the misuse of hedges for the purposes of regulatory capital arbitrage.

3. What are commenters' views on the proposed regime to strengthen the relationship between the standardized and internal models-based approaches?

There are a few points to consider here. First, we agree with the overall intent to establish a stronger relationship in the calibration of both the standardized and internal model-based approach. However, we are uncertain as to how the Committee intends to achieve this. The internal model requires banks to calibrate their data to a stressful period based on the bank's own portfolio. If the partial risk factor approach is adopted, the risk weights are calculated independently of the bank's portfolio and thus, depart from the "stressful" period notion. The fuller risk factor approach lends itself more to a consistent calibration provided that banks are required to run the model across many time periods to determine the stressful calibration. In this case, all correlations and volatilities should be taken from the stress period with no simplifying assumptions made on the intra-group correlations (right now the Committee suggests setting intra-group correlations to 0). In both cases (internal model and fuller risk factor approach) the Committee can elect to overlay their own inter-group correlations on top of the stress period calibration. Ideally, the algorithm specified in Annex C can be used to cascade the regulatory prescribed inter-group correlations down to the individual risk factors for a consistent calculation between the internal model and standardized model (based on the fuller risk factor approach).

Second, should the Committee insist on mandatory standardized measurement, then the standardized approach should be brought more in line with existing approaches that banks may already have in place in order to reduce the operational burden (and operational risk) of running separate and disjoint algorithms in parallel indefinitely. The fuller risk factor approach resembles closely the traditional parametric forms of VaR that banks may have used as part of their internal models historically. For this reason, we suggest adopting this approach.

Third, we're not sure what purpose a floor based on the standardized charge serves. Already the recommendations require banks to: (i) calibrate and calculate capital based on stressful periods, (ii) overlay regulatory specified inter-group correlations to reduce diversification benefits, (iii) calculate capital taking into account longer time horizons, (iv) move to a more conservative measure of risk (i.e., ES) and (v) employ more rigorous backtesting routines with the ability to "turn off" internal models for any given portfolio. These requirements lead to a much more conservative measure of risk which will diminish the need for additional restrictions. The Committee's observation that more complex products and risk models could benefit from an additional safeguard (Section 3.5.3) is somewhat of a redundant argument given that, should these products/models not backtest well, they will be forced to move to a fully standardized approach regardless, thus making a floor irrelevant. Our concern with the floor is not that it will lead to larger capital requirements but that, as the Committee has already pointed out, the inclusion of a floor relative to the standardized model and the intent to make standardized measurement mandatory, will almost surely distort incentives for banks to continually improve their internal risk modelling standards. The end result could be most banks adopting the

standardized model as simply a minimum compliance alternative for regulatory capital while still managing their true business operations with a different set of models and standards, thus failing the use-case.

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?

As discussed above, we agree in principle with the Committee's proposed desk-level approach and overall framework regarding the model approval process (Figure 1 in the Review). The idea of running both Attribution and Backtesting is sound in our opinion. However, the criteria[9] for determining Modellable versus Not Modellable risk factors in Step 3 is rather arbitrary and misleading, potentially leading to regulatory capital arbitrage. Frankly, having an abundance of historical data for a particular risk factor in no way ensures that a bank's model for *simulating* that risk factor is appropriate. The risk factor modelling capability should be the true test for determining whether it is Modellable or not; the historical data does not represent the whole story.

Given this, we believe the Committee's proposal can be strengthened by adopting a two-level, quantitative test. The Attribution test serves to indicate whether the desk is capturing all relevant risk factors in their analysis and is using appropriate pricing models. Backtesting includes those elements as well as risk factor modelling. In this regard Attribution is a higher level test. A bank failing Attribution would not stand much chance of passing Backtesting. A bank failing Backtesting but passing Attribution would strongly indicate problematic risk factor modelling. Thus, the Committee should use the Attribution testing as the criteria for Step 2. The Backtesting criteria should be pushed to Step 3 to determine whether the risk factor is indeed Modellable or not, thus removing the vague criteria based on historical data.

To further our argument, the Committee has suggested that Non Modellable risk factors have stress scenarios calibrated to the same confidence level as those deemed modellable (Section 4.4). However, the Committee's current criteria (i.e., a Non Modellable risk factor is one that lacks available or frequent historical data) calls into question exactly how such calibrations can occur without a distribution of data to draw from. For this reason as well, the qualitative tests currently suggested for Step 3 are ill-formed. Sparse, illiquid market data will most likely result in an Attribution failure meaning the risk will be handled in a standardized way, bypassing the need to create artificial stress scenarios.

As discussed in the Unifying Framework section above, it also would be reasonable for backtesting approaches between CCR and MR to be brought closer in line with each other. Current regulations for CCR require backtesting both risk factor and portfolio distributions. Perhaps these requirements can be mapped more consistently with the Attribution and Backtesting ideas for MR.

[9] The Committee has suggested that a risk factor's modelling eligibility be determined by evaluating the quality of the data which includes such things as availability and frequency of data.

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 Committee has claimed that the Direct method could require significant approximations (for example, filling in missing risk factors) as well as impose a significant computational burden. However, both statements are not necessarily correct. The Direct method may not require more approximations than the Indirect Method, which requires banks to arbitrarily pick their driving risk factors and then assume a constant factor between full and reduced risk factor ES figures. The comment on the computational burden is also not clear. The Committee has stated that, "*there is also a significant computational burden when searching over longer periods using all risk factors*." (Section 4.5.2) The vast majority of computational effort in a full MSMC simulation model is actually valuing the portfolios across scenarios and time steps. This burden is independent of the number of risk factors in the scenario. There really is no clear preference for either approach. In our opinion, the Committee should opt for a consistent approach with CCR, which also requires the selection of stress periods. The Committee should also emphasize consistency with the amount of historical data needed in the stress period for both CCR and MR. Current CCR requirements stipulate 3 years of historical data spanning the stress period while current MR stress period requirements stipulate 1 year.

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?

A few points to consider. First, the merits of constraining diversification benefits can be questioned, given that the internal model is already required to be calibrated to a stress period based on the bank's portfolio. Presumably this stress period, by the result of it producing the largest ES measure, will already capture the notion that diversification benefits between major risk classes may not hold under stress.

Second, the Committee's suggestion to overlay regulatory-prescribed, inter-group correlations, irrespective of a banks' actual portfolio, may not result in a conservative measure of risk after all. This is acknowledged by the Committee's use of the "max" function in Formula (1). In essence, the proposal may result in a great deal of extra computations and operational risk for banks for little end-benefit regulatory-wise.

Third, correlations can only be meaningfully prescribed at the risk factor level, since they are independent of the portfolios held. The resultant correlation between risk classes will depend on the risk factor correlations and the portfolio compositions in each risk class. The inter-group correlations specified by the Committee cannot be the same for each bank unless their portfolios have identical exposures to each risk class. Forcing a set of standard inter-group correlations on all banks is inequitable, penalizing some and benefitting others simply because their risk profile may differ in each risk class.

Fourth, the use of Formula (1) is problematic on many levels. Annex B discusses the mathematical and conceptual shortcomings with this approach and indicates clearly how it can be used to achieve regulatory capital arbitrage. It should also be noted that the Committee's suggestion for using this top-down aggregation approach is a direct contradiction to the Committee's 2009 findings, which indicated that such approaches are generally not reliable for aggregating risks:

> *"Moreover, in certain portfolios market and credit risk are related in a non-linear way. Since this means that they are inextricably linked, conventional approaches that estimate each risk type separately and then aggregate them (such as "top-down" risk aggregation approaches), which are widely used in the industry, may lead to sizable biases in overall risk estimates. For example adding the separately estimated risk components may not be conservative, as often thought, because non-linear interactions may lead to compounding effects... Claims about the presence of diversification benefits between market and credit risk should be regarded with great caution if they are not derived from an integrated ("bottom-up") approach."*[10]

While the 2009 finding was expressed from the viewpoint of market and credit risk, it is clear that general risk classes like IR, FX, EQ etc will produce compounding effects in the portfolio for many instruments and, therefore, should not be looked upon as a reliable (or even conservative) measure of aggregate risk.

Fifth, neither the desk-based nor risk-factor based approaches are desirable. The desk-based approach will rely on judgement for allocating positions to risk factor groups. Judgement opens the door to regulatory capital arbitrage, whereby a desk's allocation may be made as a result of capital savings considerations instead of reflecting the driving risk factors. As well, the desk-based approach suffers from the shortcoming of essentially ignoring major classes of risks. For instance, movements in implied volatilities are often a major contributor to portfolio fluctuations during times of stress. Instruments in all major risk factor classes (EQ, FX, IR etc...) can have exposure to implied volatilities. The Committee's current suggestion would be to wrap implied volatilities into each risk factor class (for instance EQ implied volatilities get mapped into EQ, etc.) thereby limiting the extent of direct diversification control the Committee can exercise on this important risk factor. Of course, both these deficiencies can be overcome by adopting the risk-factor based approach. However, this approach suffers from operational and computational constraints. Operationally, it may be more difficult and cumbersome for banks to split scenarios by major risk factor classes. As well, creating five major risk factor buckets would essentially mean increasing computational requirements five-fold (as the portfolio will need to be re-run under each split of scenarios).[11]

[10] "Findings on the interaction of market and credit risk". BCBS. Working Paper No. 16. May 2009. Available at http://www.bis.org/publ/bcbs_wp16.htm.

[11] Of course, banks can save on computational requirements by adopting approaches whereby they ignore positions not impacted by the particular risk factor class but the management of such complex risk-factor-to-model relationships brings about its own set of challenges from an operational standpoint.

The suggestions put forward in the Unifying Framework section above address these major shortcomings. Should the Committee continue to insist on requiring some control over diversification benefits, however, then Annex C provides a sensible and pragmatic alternative. As a first step, the Annex describes how to "roll-up" risk factor correlations to aggregate risk class correlations. Once done, the Committee can specify a set of "stress shocks" at the aggregate level (for instance a shock of +20% then -20% to all aggregate correlations). The Annex then describes how the stress shocks can be cascaded down to the individual risk factor level, producing two new "stressed" matrices corresponding to the upward and downward shocks. Banks would be required to run their internal model three times using the original matrix and the stress up and stress down matrix. As stipulated in the Unifying Framework section, regulatory capital can then be obtained as the maximum ES calculated using these three matrices. In this case, banks are only required to repeat their market risk calculations three times (instead of five times), saving computational time, and they will not need to make arbitrary decisions about how to break up risks. At the same time, regulators are still able to exert a direct influence on potential diversification benefits.

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

In general, we believe that integrated models of risk are better than a patchwork of separate capital charges. Integrated models are generally more risk sensitive and therefore less prone to regulatory capital arbitrage. Integrating market and credit risk in a transparent and computationally efficient manner is not only possible, but done routinely by many banks. Annex C in our 2008 Basel response to IRC[12] described a real-life example integrating market and credit risk based on a computationally efficient algorithm[13].

The challenges with integrating market and credit are not methodological, but rather conceptual. The current IRC charge specifies a constant-risk level liquidity horizon rebalanced to 1 year, and a 99.9% confidence interval. This was done to bring the charge in line with the Basel II credit charge, thus avoiding the potential for regulatory capital arbitrage between the trading and banking books. Moving default and migration into the proposed market risk framework will mean departing from those assumptions, potentially resulting in reduced capital levels since shorter time horizons without the requirement of rollovers is prescribed for market risk. The subsequent modelling differences for issuer risk between banking and trading book may once again entice regulatory capital arbitrage despite the Committee's intent to make the trading/banking book boundary more robust.

[12] The response can be accessed at http://www.bis.org/publ/bcbs14041/cacomments.htm.
[13] Details on the algorithm can be found in: De Prisco et al. *"Compound scenarios: An efficient framework for integrated market-credit risk"*, The Journal of Risk. Winter 2008/09 Volume 11 Number 2.

Also, it should be noted that the Committee's interest in integrating market and credit is inconsistent with other elements in the Review. For instance, the Committee's current proposal will not even allow an integrated model between standard risk factor classes like IR, EQ, FX, etc., unless the integrated model is shown to be more punitive than one based on top-down aggregation. Therefore, it would make little sense to consider tackling the concept of integrating market and credit risk when the market risk measure itself is not allowed fully integrated measurement.

As far as Attribution and Backtesting are concerned, including default and migration risk into the measure will not lead to demonstrably different mechanics for assessing either result. In this case, default and migration movements would be additional items to track and compare to theoretical results. The issue, of course, is the quantile parameter and ES measure. The more extreme the quantile, the more challenging to backtest.

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?

Both VaR and ES have their own shortcomings as measures. VaR lacks coherence as a risk measure, although this is rarely, if ever, relevant for real bank portfolios. ES is a coherent measure and takes into account more tail information than VaR, but it also has a larger estimation error (for the same number of scenarios) than VaR for some distribution types. For instance, Annex D shows that, for a normal distribution, the estimation error of ES exceeds that of VaR for quantile levels above 65%. ES will also be more difficult to backtest. Using a "number of violations" approach for backtesting may not be easily applied to ES given that it will never be clear whether a violation has actually occurred. For instance, a loss exceeding the ES amount may be completely appropriate and not indicative of any modelling problems given that ES is an expectation of all tail losses. Without knowing fully the tail's distribution, one would never be able to determine if such observations are statistically significant and therefore, a true violation of the model.

Given the above, we find the focus of which measure to use to be misguided. In the end, any measure is only as good as its underlying distribution. What failed the industry during the 2007-2008 period was not the VaR measure per se but the fact that the underlying distribution of losses being modelled was not realistic. Based on the Committee's own findings, the models failed to reflect losses actually incurred because "*many banks did not regularly update time series data and some key risk factors driving the observed losses were not incorporated into their VaR models*".[14] These elements indicate that the simulated distribution of losses was unreliable, in which case ES (apart from allocating slightly more capital) would have failed just as miserably during that time period.

Many of the Committee's suggestions in the Review are primarily directed towards improving the realism of the underlying loss distribution (i.e. introducing longer liquidity horizons, enforcing more stringent backtesting and attribution etc). Presumably such steps will serve to

[14] Analysis of VaR backtesting exceptions (Annex 1, Box 2).

improve the underlying modelling, thereby making the explicit choice of a measure on that realistic distribution somewhat irrelevant. If the distribution is reliable, then either VaR or ES should be reliable.

Taking the above points into account, we find the current backtesting procedures to be lacking. The current "stoplight" approach for backtesting violations indicates that a multiplier of anywhere between 3 and 4 should be assigned partly on the basis of how many backtesting violations occur over an interval. In our opinion, the *frequency* of violations has little to do with the multiplier (which, in our view should account for the *severity* of the loss given a violation). With this in mind, our suggestion is relatively straightforward and combines both measures of VaR and ES. Banks should continue to use VaR for regulatory capital calculations and for backtesting the frequency of violations (should the Committee not be comfortable with the lower capital figure in comparison to ES, they can always increase the quantile level for VaR). If the VaR model produces more than (say) eleven violations in a given year (assuming a 99% quantile), then the model is turned off for that trading desk and the standardized model is used[15]. If the model produces an acceptable number of violations (consistent with the quantile) then the multiplier should be based on the expectation of the ratio (Actual Loss / VaR) incurred in the violation cases. For instance, if a bank's VaR model produced three violations in a given year and the loss suffered in those three violations was on average 1.8x the VaR figure, then the model is deemed acceptable from a frequency standpoint and the bank should be allowed to use that model but with a multiplier of 1.8. In this approach, capital will be allocated based on an *ex-post* ES measure (taken from actual occurences), not one that is model driven[16]. Of course, elements concerning the number of acceptable violations and a floor for the factor can all be decided upon by regulators to ensure prudence in the final capital figure.

One potential criticism of such an approach is that it involves using 1d horizon VaR and loss figures to calculate the capital factor, when in actuality, the factor will be applied to a VaR figure calculated over potentially much longer liquidity horizons (out to 1 year). A factor calculated over shorter periods may not be relevant for longer periods. We address this issue in Annex E. The 1d factor is considered to be a conservative measure of risk when looking at normal, lognormal and jump distributions.

To recap in algorithmic form (reflecting the points already made in our response to Question 4), the decision criteria in Figure 1 of the Review should be thought of as:

Step 2
If P&L attribution at trading desk level fails, then use Standardized approach for trading desk
Else, proceed to Step 3

[15] A likelihood ratio test set at a 1% significance level indicates that a 99% VaR model will be rejected if it produces 11 exceptions or more during a 1 year period.

[16] As a point of reference, we analyzed the financial statements of various banks over the last 5 years that disclosed the number and severity of their VaR exceptions. Filtering out rogue trading losses, the remaining violations resulted in average factors ranging from 1.1 to 2.5. In comparison, for a standard normal distribution, the ratio of ES and VaR at a 99% quantile level is 1.145.

Step 3
Perform backtesting.
If frequency of violations exceeds a threshold, then risk factor / desk is considered not modellable and move to Stress Scenario approach.
Else, use internal VaR model but with ex-post ES adjustment for calculating factor.

9. Which of the two approaches better meets the Committee's objectives for a revised standardised approach?
10. Do commentors propose any amendments to these approaches?

It is difficult to imagine how the partial risk factor approach satisfies many of the Committee's requirements for a standardized model. The calibration of RWs is hardly clear or logical and will become much more complex than for the simple EQ example the committee has put forward in Annex 6. Also, the approach is far from simple and transparent, as by the Committee's own description, "special" rules will need to be developed to account for:

- Instruments without market values
- Instruments with cross-cutting factors
- Instruments with delta equivalents and delta hedges
- Instruments requiring general interest rate risk treatment
- etc...

Given that the proposal will allow banks to apply or "turn-off" internal models per trading desk, it seems reasonable that those trading desks having the most exotic products will run the greatest likelihood of failing internal model tests and, therefore, resort to a standardized approach. The exotics (whether they be credit related or something else) could prove to require an endless array of additional rules to capture their risks in a meaningful way in the partial risk factor approach. Without a string of customized and unwieldy mapping rules to cover exotics, the partial risk factor approach will fail precisely where it is needed most. In this regard, it hardly appears to function well as a credible fallback.

The fuller risk factor approach is preferred as it resembles very closely parametric models for VaR which many banks are already familiar with and which several banks have used in the past as an internal model. These models are generally well understood in the industry and can capture many of the elements the Committee seeks. However, there are two slight amendments that would be advisable to this approach.

First, one suggestion would be to employ the general "tweaking" approach everywhere, for all instruments regardless of complexity in order to determine sensitivities and exposures to various risk factors. This bypasses the need to create arbitrary rules for linear instruments. In our experience, fewer arbitrary rules lessen the likelihood of operational difficulties in running the system. Tweaking everywhere will allow a rigorous and consistent approach for generating the required risk factor mapping.

Second, while we understand the purpose of the risk factor hierarchy, the Committee needs to be very careful in the specification of volatilities to avoid double-counting risks up through the hierarchy. For instance, in the example provided in Annex 7, Daimler would have a beta to the Industry Index and the Industry Index would have a beta to the World Index. If the Committee used the full volatilities for both the Daimler stock and Industry index, they would be triple-counting the risk of Daimler stock up through the hierarchy. More specifically, Daimler's total variance (i.e. the variance of its stock price) can be represented as:

$$\sigma_D^2 = \beta_{DI}^2 \sigma_I^2 + \sigma_d^2 \qquad (1)$$

where
σ_D^2 , represents the total variance of Daimler stock price
β_{DI}, represents the beta of Daimler with the Industry index
σ_I^2, represents the total variance of the Industry index
σ_d^2, represents the idiosyncratic variance of Daimler

Likewise, the Industry index's variance can be represented as:

$$\sigma_I^2 = \beta_{IW}^2 \sigma_W^2 + \sigma_i^2 \qquad (2)$$

where
β_{IW}, represents the beta of the Industry index with the World index
σ_W^2, represents the total variance of the World index
σ_i^2, represents the idiosyncratic variance of the Industry index

Substituting (2) into (1) yields:

$$\sigma_D^2 = \beta_{DI}^2 \beta_{IW}^2 \sigma_W^2 + \beta_{DI}^2 \sigma_i^2 + \sigma_d^2 \quad (3)$$

According to (3), even if all betas are set to 1 (which is what the Committee has implied doing), in order to not double-count risks for Daimler, the full variance (volatility) of the Worldwide equity should be used but only the idiosyncratic variances of the Industry index and Daimler stock should be used.

It is also unclear where the Committee intends to get volatilities for the Level I factors as these types of aggregate, worldwide indices may not exist in reality. For these reasons, it may be easier to consider a 2-level hierarchy only (Levels II and III in the proposal). While it limits hedging at the highest levels, it also avoids the extra volatility term in Level I. It would also avoid much of the double-counting of risks as a result of a hierarchy that is too extensive.

To summarize, it is our view that the Unifying Framework described above not only brings much needed consistency between MR and CCR but also helps to address many of the shortcomings present in many of the proposals put forward in the Review. Re-using many of the CCR modelling elements for MR is, in fact, the ideal integrated framework the Committee

should be striving for. While we agree, in principle, with the Committee's overall intent to build an integrated framework via the 3 step outline in Figure 1, **the introduction of many of the modelling restrictions/constraints in Step 3 only serve to make the final framework less risk-sensitive and more prone to opportunities for regulatory capital arbitrage**. Despite some of the more technical issues, we do applaud the Committee on its overall initiative to remedy the challenges brought about by the current patchwork of regulatory measures.

We thank the Committee for its diligent review of our concerns and comments. We welcome queries or requests for further detail on any of the topics raised, or related issues. We can be reached via email at Ben.Deprisco@ca.ibm.com or telephone at 416-217-1500.

Sincerely,

Ben De Prisco
Head of Research and Financial Engineering

Overview

The Committee has proposed the use of varying liquidity horizons for different positions in the trading book portfolio, recognizing the fact that there is a range of liquidity observed in traded instruments during periods of stress. Regulatory capital should be allocated based on the perceived ability to "liquidate" the instrument at the end of its particular liquidity horizon. The concept of liquidation bypasses the need to model complex rollovers in the portfolio (unlike IRC). However, the Committee has highlighted 2 particular cases which, in their opinion, will require rollovers:

- When instruments mature before the end of their liquidity horizon; and
- When hedges mature before the end of the hedged instrument's liquidity horizon

To overcome these isssues, the Committee has suggested the use of "Instantaneous Shock" scenarios (section 4.5.5) and has discussed various options for applying these types of shocks with scaling factors (Annex 4, Section 2). The purpose of this Annex is to show that the Committee's rationale for rollover modelling is conceptually flawed and that the recommended use of instantaneous scenarios will result in an inappropriate allocation of regulatory capital, leading to clear cases of regulatory capital arbitrage.

Conceptual Flaws

We disagree with the Committee's first point - instruments expiring before the end of their liquidity horizon, in fact, WILL have regulatory capital assigned. The Committee is under the mistaken impression that the only simulation choices available are either instantaneous shocks or single-step shocks placed at the end of the liquidity horizon (as discussed in their 3 options in Annex 4, Section 2). However, as proposed in the Unifying Framework section in our response to the Committee, the adoption of an MSMC framework, which is already supported by the Committee for CCR, can be used to capture risk at the maturity of the instrument. By generating a series of risk factor paths at multiple timesteps throughout the horizon (each date being considered a risk factor "trigger date"), the bank will be able to capture the dispersion in instrument values at its maturity date. To be more precise, the risk and volatility will be captured up to the closest trigger date preceding the instrument's maturity. As an example, if we assume a particular instrument has a 6 month liquidity horizon but only a 3.5 month maturity and the bank has adopted a MSMC simulation using monthly trigger dates, then the "risk" of the instrument will be based on 3 months of volatility. Closely-spaced triggers (as is the case for CCR) will always ensure a meaningful risk allocation for any expiring instrument. Examples 1-4 below illustrate this idea and the overall MSMC approach in more detail.

The second point requires further elaboration by the Committee as either: (i) we have misunderstood the point or (ii) the Committee is contradicting itself. The Committee has defined a liquidity horizon as *"the time required to sell a financial instrument, or hedge all its material risks, in a stressed market, without materially affecting market prices."* (section 3.3.1) The

Committee then goes on to elaborate that, *"a liquidity horizon of, say, three months would mean that the calculation of the regulatory capital charge would assume that the bank can hedge or exit its risk positions* ***after*** *three months and not require any rebalancing assumptions."* (section 3.3.2, our emphasis) If a hedge matures before the end of the hedged instrument's liquidity horizon, then **by the Committee's own definition** of liquidity horizon, the hedge cannot be rolled over within the liquidity horizon. If it could, then the underlying exposure should have been granted a shorter liquidity horizon, by definition. So either: (i) there is confidence that the hedge can be seamlessly rolled over in which case, the underlying exposure picks up the hedge's maturity date as its liquidity horizon, or (ii) there is little confidence that the hedge can be rolled over, in which case, the exposure is deemed "naked" from the hedge maturity to the liquidity horizon. In either case, **there is no need to model complex rollovers of the hedge**.

The Committee may argue against the interpretation above by claiming that the definition of liquidity horizon implies 'without materially affecting market prices', thus allowing banks to model the hedge rollover within the liquidity horizon but with some slippage in price and effectiveness. However, such an argument flies contrary to the Committee's own guiding principle that *"the capital framework should only recognize hedges if they are likely to prove effective - and can be maintained - during periods of market stress."* (section 3.4) Allowing the modelling of slipping hedges would be so arbitrary as to lack prudence and conservatism in the overall framework and run contrary to the Committee's guiding principles.

Based on the points above, we see no requirement to actually model rollovers for the calculation of regulatory capital. As such, there would be no justification for requiring the use of "instantaneous shock" scenarios. As discussed below, such approaches de-sensitize the risk profile, leading to opportunities for capital arbitrage.

Issues with Instantaneous Shock Scenarios

The Committee has already fully acknowledged that use of instantaneous shock scenarios does not capture the risks of exotic instruments in detail. *"For example, no particular time paths are specified for the risk factors, which means that the risks from path-dependent instruments (eg Asian options, barrier options) are not well captured."* (Annex 7, Step 3) While this point is clear, what is often overlooked is that even the simplest and most common trading book instruments are path dependent. For instance, swaps, interest rate caps/floors and floating bonds all have resetting rates which can impact risk profiles immensely. The Committee also overlooks the notion of time decay and pull-to-par. For instance, zero coupon bonds, options, swaptions, etc., all have significant amounts of time decay or pull-to-par which will impact capital allocation.[17] These elements can only be captured properly through an MSMC simulation.

[17] As an example of pull-to-par, the risk associated with a zero coupon bond maturing in 2 weeks but having a liquidity horizon of 1 month should be zero (ignoring credit risk) as the value of the instrument at maturity is known and market independent. Instantaneous shock scenarios, however, would show risk for this instrument.

In the various examples that follow, we describe at a high level the notion of MSMC simulation, the benefits it can bring to the calculation of market risk and how it can prevent regulatory arbitrage (as compared to Instantaneous Shock scenarios).

Example 1 - Single Stock

Consider a portfolio consisting of a single stock. We can simulate the risk of that stock through time assuming the underlying process is a Geometric Brownian Motion (GBM)[18]:

$$S_\tau = S_0 e^{\left(\beta - \frac{\sigma^2}{2}\right)\tau + \sigma W_\tau} \tag{1}$$

where
S_0= current stock value = 6,761
S_τ= stock value at any future trigger time τ (where τ is expressed in units of a year)
β = drift on the stock = 0
σ = annual stock volatility = 30%
W_τ= standard Brownian motion

Simulating 10k scenarios with trigger times: 2w, 1m, 3m, 6m and 1y produces the following graph:

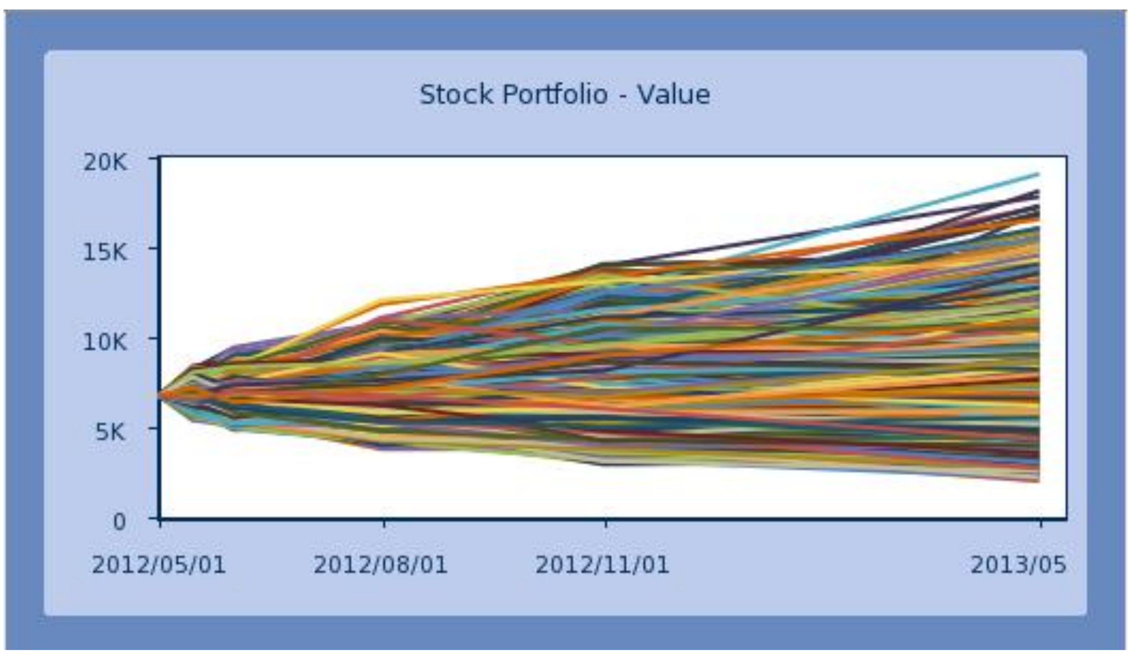

Figure 1: Stock profile through time

[18] Note that the use of GBM is just for illustration purposes and not a requirement for implementation of the proposed Unifying Framework. The process for S_τ is written under a historical measure, not a risk-neutral measure. As such, there is no requirement for β to be set to a risk-neutral drift.

In the example above, the MSMC simulation generates a set of paths for the value of the equity with the resultant distribution of equity values growing through time (as would be expected in a GBM process). In this example, both the simulation of the stock risk factor and the holding of the stock go out to 1 year. The final distribution at 1y produces the following results: ES(99%,1y horizon) = 3,892 and VaR (99%, 1y horizon) = 3,600. The ES allocates 8% more capital than VaR in this case.

Example 2 - Single Stock with 3m Liquidation

The example above had the simulation of the stock risk factor and the holding of the stock synchronized to the 1y liquidity horizon. However, there is nothing that forces the two being synchronized. In fact, the major benefit of the Unifying Framework is that the underlying risk factor simulations can be done independently of the instrument's valuation and liquidity horizon. For instance, if we assume the stock's liquidity horizon is now 3m instead of 1y, we can re-use the same risk factor paths from Example 1 but this time model a liquidation at the 3m point. The "liquidation" has the effect of dropping the stock into cash at the 3m point. The cash amount then is carried forward at a constant value to the end of the 1y period.[19] This process produces the profile below:

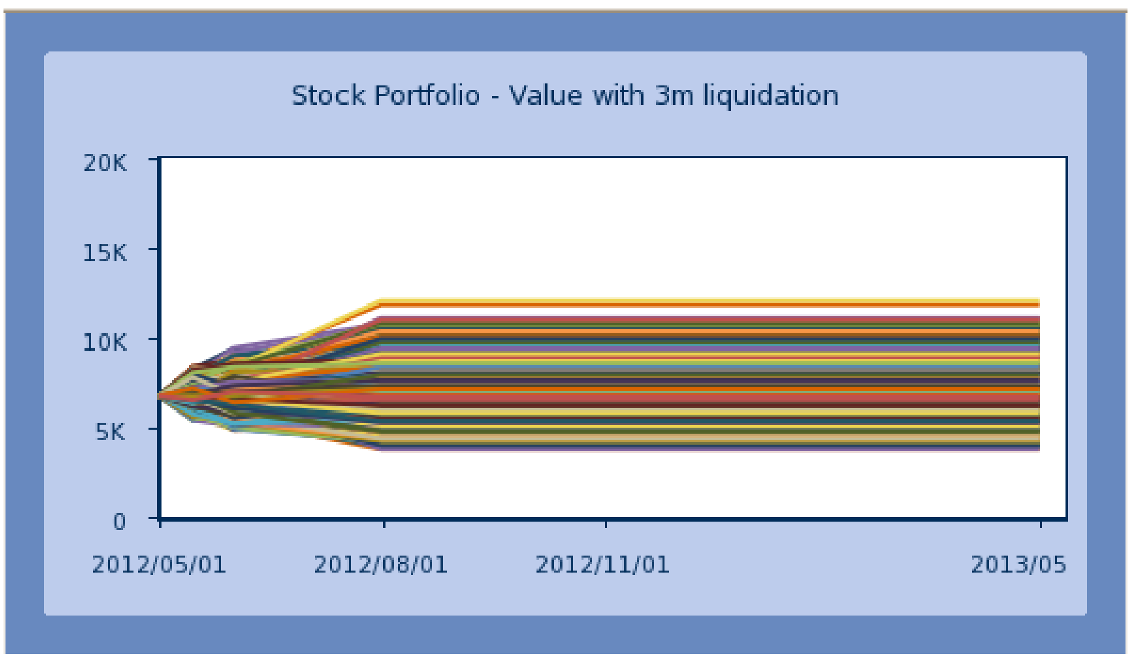

Figure 2: Stock profile assuming 3m liquidation

[19] Note, assuming a constant value for cash is one of several assumptions that can be made. Cash (both positive and negative balances) could also be assumed to grow by a given rate (which can be either stochastic or deterministic) to the 1y point.

Figure 2 shows the stock's volatility "stabilizing" after the 3m point due to the "liquidation" of the holding into cash. For this example, calculating the risk measures at 1y or 3m would be equivalent (since there is no further source of variability beyond 3m). In this case, the risk measures are ES(99%,1y horizon) = 2,348 and VaR (99%, 1y horizon) = 2,122. Reducing the liquidity horizon to 3m shrinks the ES by 40% and the VaR by 30% in comparison to 1y. More importantly, this simple example should emphasize to the Committee that the discussion of instantaneous scenarios with scaling factors is unnecessary when trying to combine varying liquidity profiles together in one metric. The Unifying Framework offers the following general and robust solution for combining liquidity horizons:

1. All risk factors should be simulated together on all trigger dates, thus developing a full set of risk factor paths to the end of the 1y horizon. These scenarios will not concern themselves with liquidity horizons but rather, with trigger dates. The consistent simulation of all risk factors will preserve their implicit correlations through time (i.e. there is no need to explicitly correlate distributions at different time steps).
2. Liquidity horizons are assigned to each instrument. At its liquidity horizon, each instrument will be settled into cash, thereby "locking in" its profit and loss distribution at that point.
3. Regardless of liquidity horizons, the distribution at the 1y point is the only distribution that matters. It will serve as a proper amalgamation of all instruments and all liquidity horizons. For instance, adding a second stock with a 6m liquidity horizon would result in the profile having additional volatility until the 6m point when the second stock is liquidated. Just as before, the 1y distribution fully reflects the different liquidity horizons as well as the correlations between stocks. For instance, if the two stocks were highly positively correlated, you could expect a path that produces a large gain on the first stock at the 3m horizon to also produce a gain on the second stock for the 6m horizon and vice versa.

Example 3 - Single Stock with 3m Liquidation and Illiquidity Jump

We agree with the general principle that certain instruments should have a capital add-on for jumps in their liquidity premia and that the capital charges should be calibrated based on the price experience of similar instruments in previous stress periods. However, these charges should only be treated as a standardized capital charge when using the standardized model. For the internal model, the Unifying Framework offers a more direct and intuitive way to apply the charge. In this case, the charge can be applied directly to the liquidation of the instrument using a Bid-Ask spread (BAS). For instance, revisiting Example 2, a regulator might stipulate that the stock is so prone to jumps in illiquidity that a 10% charge would be advisable. Re-running the example above with a 10% BAS has the effect of dropping the value distribution 10% lower upon liquidation[20]:

[20] Note we describe the charge as BAS and not simply a Hair Cut to valuation since liquidating a negative position should result in paying a premium on the price, not a discount.

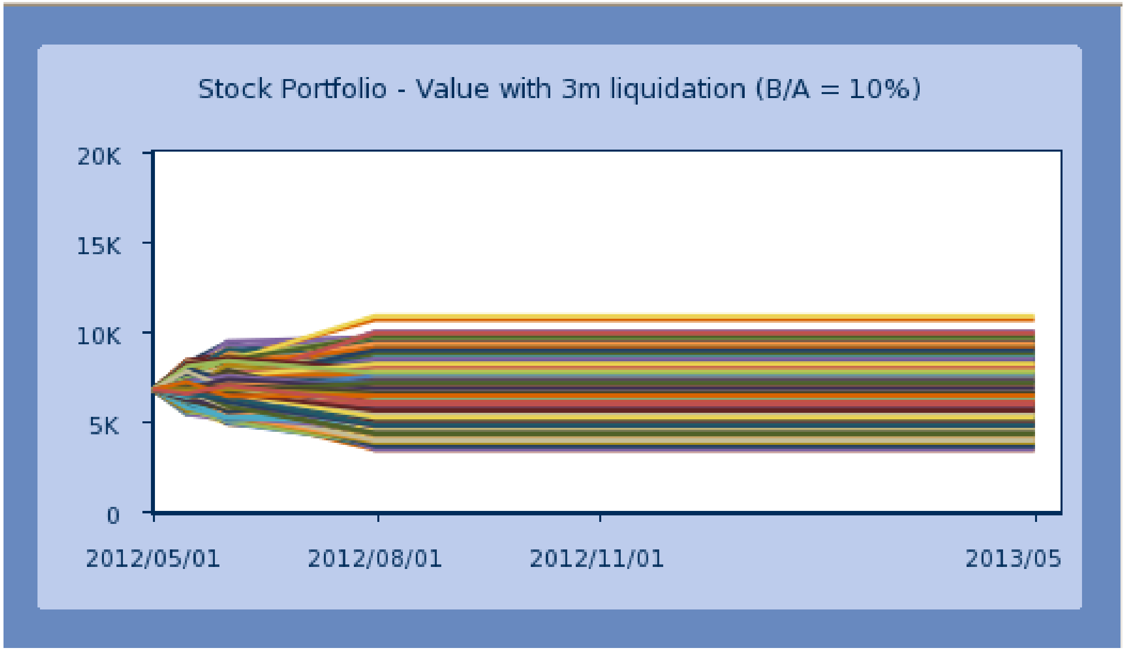

Figure 3: Stock profile assuming 3m liquidation and 10% BAS

In this case, the ES(99%, 1y horizon) rises to 2,789 and the VaR(99%, 1y horizon) rises to 2,586. In this example, a 10% BAS illiquidity charge translates to 19% higher ES and 22% higher VaR.

It should be noted that if all instruments in the portfolio share the same BAS illiquidity adjustment, then the adjustment to the final ES or VaR risk measure can be calculated in a straightforward manner as:

$$Measure_{BASAdj} = Measure_{orig} + BAS(Value_0 - Measure_{orig}) \qquad (2)$$

In reality, different charges will apply to different instruments thereby making formula (2) much more complicated but the risk measure can easily be solved in the Unifying Framework through simulation. This suggestion will greatly simplify the regulatory calibration and supervision of the charge. Regulators would simply need to calibrate the BAS charges based on historical price experience during stress periods and relate those charges to banks as BAS adjustments to particular instrument profiles.

Example 4 - Option Expiring before Liquidation Horizon

This example attempts to address the Committee's misconception that an instrument expiring before its liquidity horizon would end up with zero regulatory capital. As discussed earlier, the Unifying Framework, which relies on MSMC simulations, would allow risk to be captured up to the maturity of the instrument. In this case we consider a short position in a European Call Option on the same stock as in the previous examples. We re-use the stock paths from the previous examples and at each time step and scenario we re-value the option using the Black-Scholes formula with an Implied Volatility = 30%, Strike Price = 6,750 (making the option at-

the-money), a maturity of slightly over 1m and a liquidity horizon of 3m.[21] The current value for the option is -249 (given its short position) and the resultant simulation profile looks as follows:

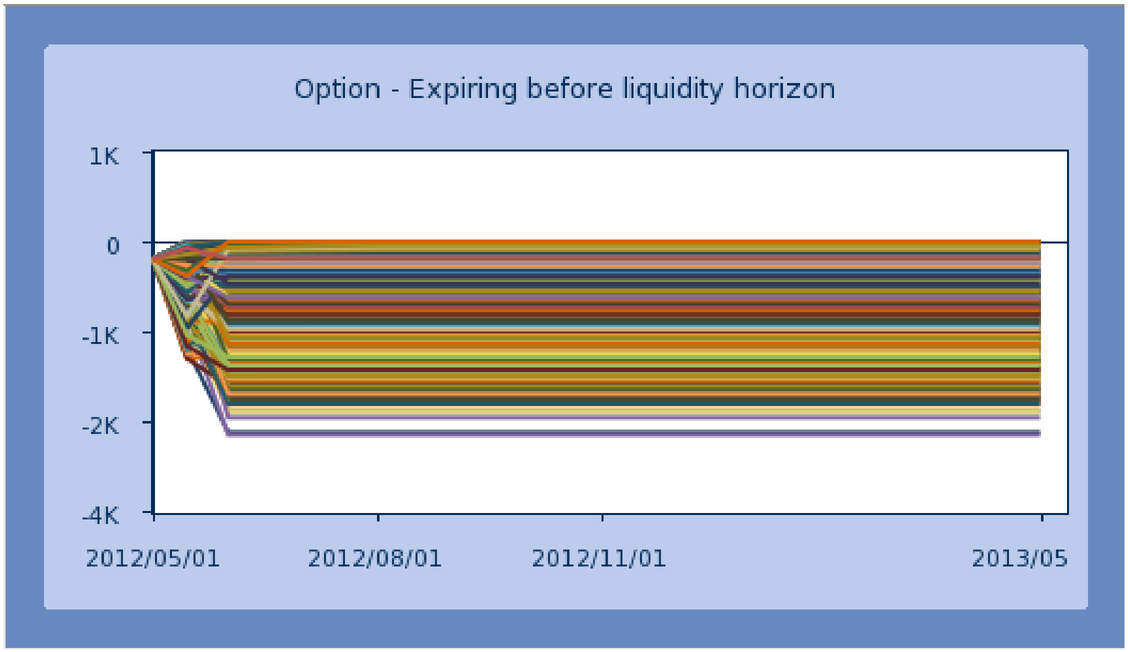

Figure 4: Option profile expiring before its liquidity horizon

Figure 4 indicates that rather than having the instrument "liquidate" into cash at its 3m liquidity horizon, it "settles" its value distribution into cash at its maturity date (in this case, the 1m point in simulation). As with liquidation, the maturity settlement results in the instrument value distribution being locked into cash for the remainder of the horizon. There are two key differences between liquidation and settlement in the Unifying Framework:

1. A liquidation is a "forced" move. It is independent of the terms and conditions of the instrument per se except for its liquidity horizon (which dictates the timing of the forced move). In contrast, settlement applies to an instrument based on its simulation profile through time and its terms and conditions. For instance, the following items would require settlement into cash:
    a. Cash flows generated from the instrument
    b. Occurrence of knock-out conditions
    c. Early exercise for American options
    d. Maturity date

    These items would impact the final profit and loss distribution for the instrument at the liquidity horizon. For instance, the regulatory capital for a bond should not only be a function of the present value distribution of the bond's remaining cashflows at the liquidity horizon (i.e. liquidation values) but should also include the cashflows (i.e. settlement) that have been collected on the bond up to the liquidity horizon.

[21] Note in this case we assume the implied volatility and interest rates are static to simplify the example and analysis. However, they are fully simulatable in the full model.

2. The BAS adjustment for market illiquidity only applies to liquidation and not to settlement.

Based on the simulation above, the following results are taken from the 1y distribution: ES(99%,1y horizon) = 1,649 and VaR (99%, 1y horizon) = 1,402. Clearly, the risk is not zero, as assumed by the Committee.

Example 5 - Option Expiring before Liquidation Horizon using Instantaneous Shock Scenarios

As discussed in the "Issues with Instantaneous Shock Scenarios" section, instantaneous scenarios do not properly capture the risk profile of any path-dependant instrument, whether it be exotic or common. Hence, they do not provide a reliable assessment of risk. **Moreover, instantaneous scenarios will not be a reliable indicator of risk for any instrument expiring before its liquidity horizon**. To illustrate and emphasize this point, consider again the option in Example 4. This time, however, instead of MSMC simulation, we simply use formula (1) to derive a set of 3m shocks to the underlying stock (consistent with the option's liquidity horizon), and apply these as a set of instantaneous shocks (IMC) in the revaluation of the option. Expressed in this manner, the Committee should see that the suggested use of IMC is flawed. If the 3m scenarios are to be believed (i.e. they are considered to be an accurate representation of how much the stock can move in a 3m period) then it is inappropriate to apply them to an instrument that expires in a much shorter period (i.e. 1m). To do so would imply much more volatility to the instrument than would be possible. This will most surely lead to regulatory capital not reflecting risks being taken and therefore increase opportunities for regulatory capital arbitrage.

Returning to the example, applying the shocks instantaneously produces the following results: ES(99%, Instant) = 3,100 and VaR(99%, Instant) = 2,637. In comparison to the results in Example 4, applying instantaneous shocks produces an unreliable increase in the risk of 88%. In order to understand how significant the differences between MSMC and IMC can be, we re-ran Examples 4 and 5 varying the volatility of the stock (σ) in our scenarios as well as the option maturity (t).

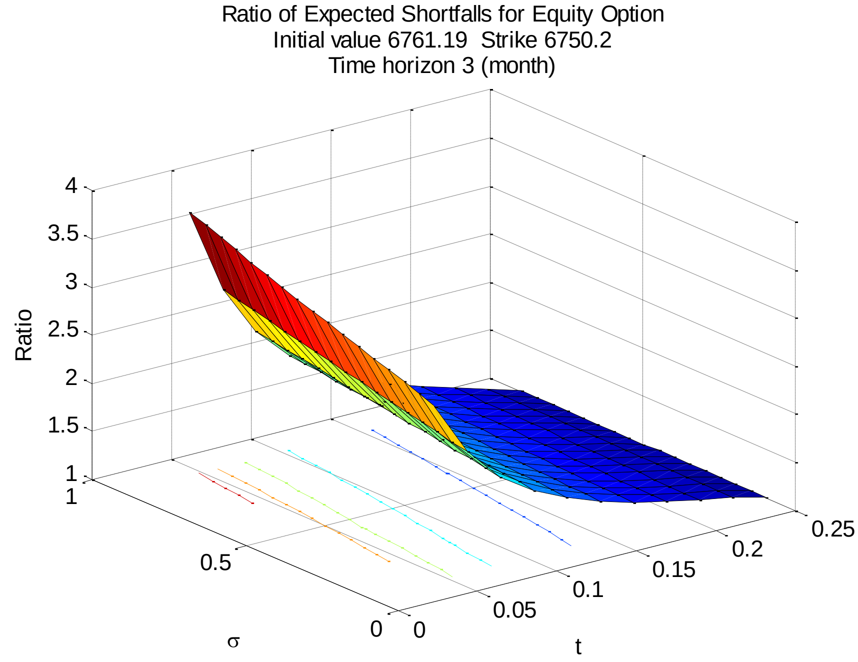

Figure 5: Ratio of ES for at-the-money option

Figure 5 plots the ratio ES(IMC)/ES(MSMC) for a 99% confidence level across a range of stock volatilities and option maturity dates. As expected, when the option expires at the liquidity horizon (t=0.25), there is little difference between IMC and MSMC (ratio close to 1). However, the shorter the maturity of the option, the greater the mis-allocation of capital. For short-dated instruments, the differences can exceed 3.6x the appropriate level of capital.

Re-running the simulations with an in-the-money option produces a very similar profile:

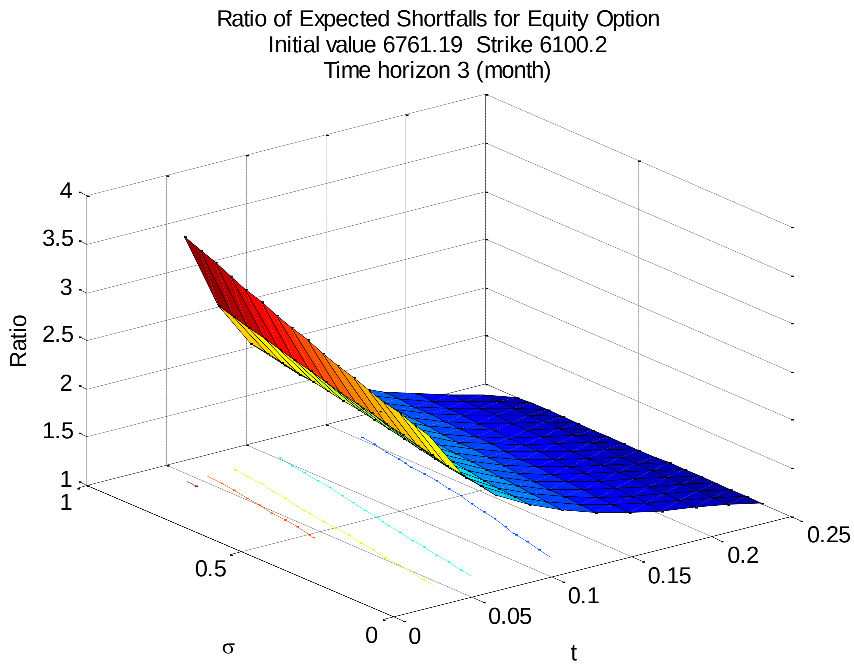

Figure 6: Ratio of ES for in-the-money option

Re-running again with an out-of-the-money option produces a similar result with slightly more dependence on the stock's volatility:

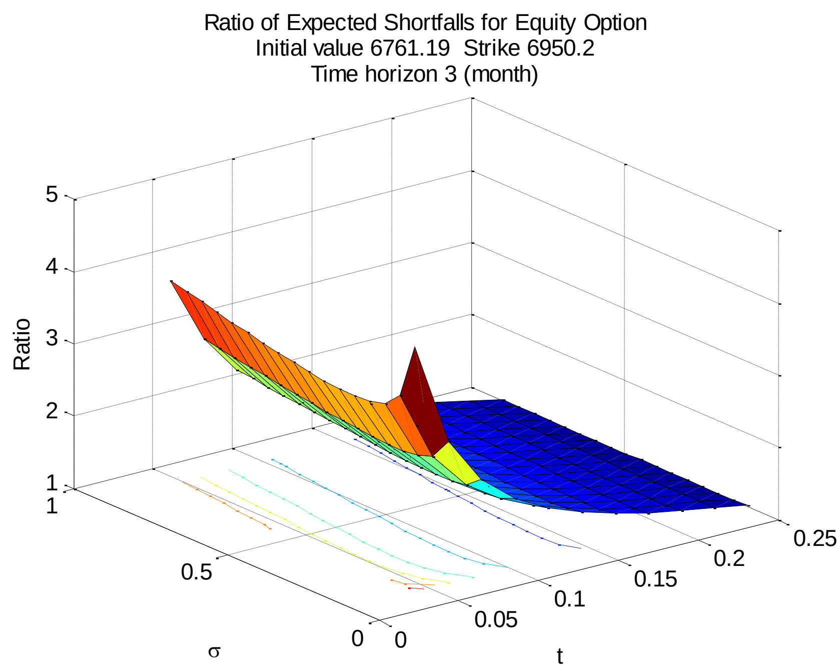

Figure 7: Ratio of ES for out-of-money option

Example 6 - Hedges Expiring Before Liquidity Horizon

As discussed in the "Conceptual Flaws" section at the beginning of this annex, hedges expiring before the liquidity horizon of the underlying exposure should not be rolled over. Doing so would be contradictory to the Committee's own definition of liquidity horizon. Either the hedge rollover is deemed to be stable and the exposure's liquidity horizon should then shrink to the maturity of the current hedge, or the hedge rollover is considered to be unreliable and therefore, the exposure should remain 'naked' for the period from the hedge maturity to the liquidity horizon.

In what follows, we construct a simple example to show that the IMC method is completely unable to capture the risks properly when a hedge expires before the exposure's liquidity horizon assuming the position is considered to be 'naked' for the time between hedge maturity and liquidity horizon.

Consider a portfolio containing 100 units of a long stock (the exposure) and a short forward contract on that stock (the hedge). The stock is assumed to have a liquidity horizon of 3m and the forward contract expires in 2m[22]. The stock's price is 1,405 and its volatility is 40% annually. As before, we use the formula in (1) to generate a series of paths for the stock price. We simulate the stock's value at 2w and then each month out to 1 year. As before, we use 10k scenarios both in MSMC and IMC modes (for IMC, we generate scenarios based on the 3m liquidity horizon). Not surprisingly, the IMC approach returns 0 risk for both ES and VaR measures. This is a result of the fact that the stock has a delta of 1 and the forward has a delta of -1 when looking at things from an instantaneous perspective.

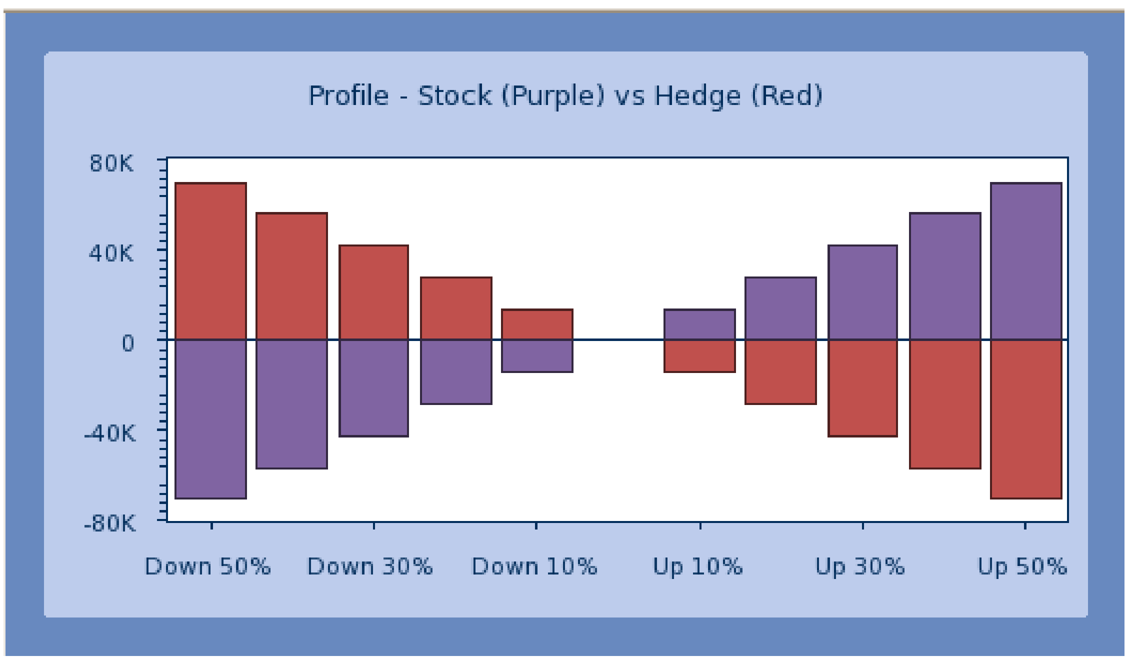

Figure 8: Instantaneous shocks with stock and hedge

[22] The forward is assumed to be cash-settled upon maturity.

Figure 8 shows quite clearly that for any instantaneous shock on the stock's value (x-axis) the forward contract hedges the stock's exposure perfectly. This is the same manifestation of non-sensical behaviour we highlighted in Example 5 where a short-dated instrument was deemed to have much more volatility than possibly dictated by its maturity date (in this case, it is the hedge which behaves inappropriately).

Simulating this portfolio in the Unifying Framework with MSMC yields a more realistic profile of risk through time:

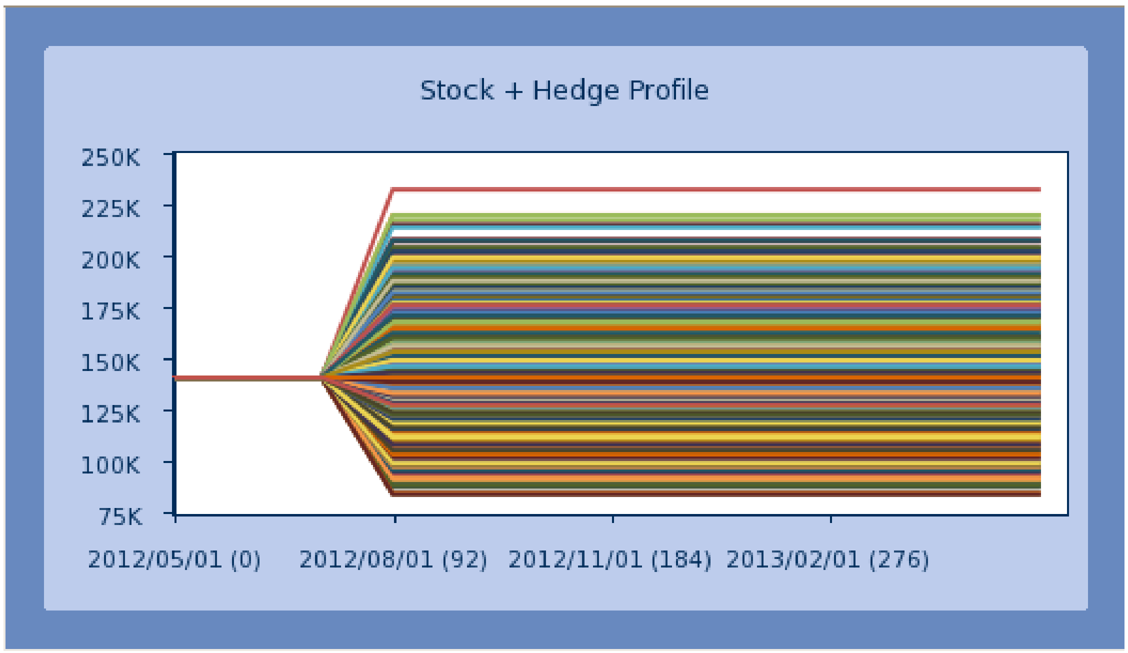

Figure 9: Profile of stock and hedge

As would be expected, Figure 9 shows that for the first 2 months of simulation, there is a perfect hedge in place for the stock. Once the hedge expires however, the naked stock has effectively 1 month worth of risk until its liquidity horizon passes. The Unifying Framework properly accounts for this risk: ES(99%,1y horizon) = 42,252 and VaR (99%, 1y horizon) = 36,878.[23] As expected, the risk grows as the distance between the hedge maturity and liquidity horizon grows:

Hedge Maturity	2w	1m	2m
ES (99%)	59,713	55,633	42,252
VaR (99%)	53,064	48,953	36,878

Table 2: Risk measures per hedge maturity

As Table 2 indicates, the shorter the hedge, the longer the period the stock is 'naked' and therefore the greater the risk. Meanwhile IMC scenarios would show zero risk everywhere. Taken to its logical conclusion, a bank could enter into a 1d forward contract on the stock. This

[23] While it may be tempting for the Committee to suggest that the use of IMC scenarios can still be used reflecting the 1m period where the stock is 'naked', the idea is short-lived when you consider that you'd have to create customized horizon scenarios for each exposure net of its hedge and then somehow amalgamate them all together.

forward contract would really not alter the risks of the portfolio at all, however it would lead to a perfect hedge and zero risk allocation when using IMC scenarios. Clearly such opportunities for regulatory arbitrage should not be permitted. For this reason, the Committee should not allow the use of instantaneous shock scenarios. Only the Unifying Framework, through the use of MSMC scenarios, can properly capture the risks in the portfolio.

Summary

The Committee's rationale for suggesting the use of rollovers is conceptually flawed. Moreover, the recommended use of instantaneous scenarios as a suitable work-around for modelling such rollovers will provide opportunities for regulatory capital arbitrage. Contrary to the Committee's belief, instruments expiring before their liquidity horizons will show appropriate capital requirements when using MSMC techniques. As well, modelling hedge rollovers before the end of the liquidity horizon is clearly contradictory to the Committee's own definition of liquidity horizons. Examples 1-4 in this annex explain the mechanics of the Unifying Framework and how it can be used to fully model trading book risks in a reliable way including jumps for illiquidity. Examples 5 and 6 point out the inappropriateness of Instantaneous Shock scenarios and how they can be used to promote regulatory capital arbitrage in a very real way.

Overview

The Committee has proposed that banks calculate separately the risks associated with major risk classes (Interest Rates, Foreign Exchange, Equities, Credit Spreads, Commodities) and then aggregate these risks using a top-down aggregation approach with supervisor-specified correlations between classes. The required capital would then be the maximum of the capital under the regulatory specified aggregation approach and that of the fully integrated model. The motivation is that this will constrain the diversification benefits banks can recognize. Our response to question 6 in the letter has already raised the following concerns with this approach:

- The top-down aggregation may be deemed redundant as banks are already required to calibrate their internal model to stress periods which, by virtue of having produced the greatest losses, should already reflect reduced diversification benefits between risk classes.
- The assignment of regulator-specified correlations does not take into account the bank's actual portfolio so it is unclear whether the aggregation will actually be a conservative figure or not (hence the need for a max function in the Committee's aggregation scheme).
- Forcing the same absolute levels of inter-group correlation on all banks will penalize some and benefit others depending on their exposure makeup in each risk class. This would be inequitable from a capital perspective.
- The use of a top-down aggregation approach is contrary to the Committee's earlier findings that such approaches are not reliable indicators of risk.
- The application of such an approach will require judgement when mapping hybrid instruments to a risk class. The presence of judgement will allow for the presence of regulatory capital arbitrage (i.e. allocations are made based on what produces the lowest capital, not necessarily what most drives the risk).
- The approach misses major risk factors such as Implied Volatilities (which would be mapped into the five categories above), reducing the regulator's ability to control diversification benefits.
- A risk factor splicing approach (i.e. splitting scenarios by risk classes) can be used to bypass the need to allocate positions to risk classes but such approaches come with operational and computational performance challenges.

The Committee may summarily dismiss these points as being irrelevant given the presence of the Max function in formula (1) in the Review. If the aggregation function produces a greater risk number then it is used, otherwise the integrated model is used. So even if the aggregation can be manipulated to produce a lower number than the integrated model (which is quite likely as we will show below), regulators will only rely on it in the event it produces a greater number. So there is apparently little downside for the Committee to suggest using formula (1).

However, there are REAL consequences for banks in terms of additional operational risks and additional calculation challenges in finding appropriate ways to split up risks and incorporating

these approaches in a robust workflow daily. The Committee should not downplay such operational risks and challenges. Also, as suggested above, the Committee has already found that such top-down aggregation approaches are not reliable indicators of risk. Sanctioning their use now potentially opens the door for banks to use them in other areas where the Committee may be less enthusiastic (i.e. there would be no logical reason why banks could not adopt such approaches for their internal model as well). To this point, there would be little credibility to the Committee's statement that *"in the context of ES, the Committee is of the view that any methods which do not rely on full repricing would not be appropriate given the importance of modelling the tail of the loss distribution."* (Section 4.5.1) What should be appropriate in one case, it can be argued, should be appropriate in the other.

Matrix Validity

Another concern relates to the Committee's vagueness about how the inter-group correlations will be specified and whether each regulator will be required to use the same estimates or their own estimates - a potential for regional regulatory arbitrage.

Given 5 general risk classes, there will be the need to specify 10 inter-group correlations. Clearly these figures must be carefully specified to ensure that the resultant 5x5 regulatory correlation matrix is a valid one.

The Committee has also suggested the use of parameter floors (or ceilings) in the specification and has even suggested that the parameters themselves would have to reflect whether the bank was long or short the specific risk factor (see footnote 37 in the Review). We find such discussions puzzling. The correlation between two risk factors should remain the same regardless of being long or short the risk factors. The "L" in Formula (1) will reflect a long or short position. Trying to specify a different set of correlations or correlation parameters conditional on the "L" values is perhaps the quickest way to ensure the matrix, and any resultant calculation using the matrix, is invalid and non-sensical.

Example 1 - Compounding and Diversification Effects

As mentioned in our response to question 6 in the letter, the Committee has already openly acknowledged that the use of top-down aggregation approaches are not reliable for estimating either compounding or diversification effects between risk factors. These effects are hardly esoteric; they can exist in even the simplest portfolios.

For instance, consider a portfolio consisting of a single stock denominated in a foreign currency. The portfolio has both EQ and FX risk, which interact in a multiplicative manner. We can simulate the stock price and the fx rate through time assuming each underlying process is a Geometric Brownian Motion (GBM)[24]:

[24] See footnote 18 in Annex A.

$$S_t = S_0 e^{\left(\beta_s - \frac{\sigma_s^2}{2}\right)t + \sigma_s W_t^s} \quad (1)$$

$$F_t = F_0 e^{\left(\beta_F - \frac{\sigma_F^2}{2}\right)t + \sigma_F W_t^F} \quad (2)$$

where
S_0= current stock value = 100
F_0= current fx rate = 1.4
β_S = drift on the stock = 5%
β_F= drift on the fx rate = 0%
σ_S = annual stock volatility = 50%
σ_F = annual fx volatility = 50%
W_t^S and W_t^F are standard Brownian motions correlated with ρ = 0.7
t = time horizon = 3m

The positive correlation between the stock and fx rate produces a compounding effect. For instance, as the fx rate falls, the stock also falls, resulting in a compounding of losses in this case. Simulating this portfolio with 10k scenarios and a 3m horizon produces an integrated risk measure of ES(99%, 3m) = 101.44. A tail decomposition analysis indicates that both the stock and fx factor contribute equally to the tail loss (i.e. they both contribute 50% of the total loss).[25] In this case, there is no clear dominant risk factor meaning that the Committee's second approach, which maps portfolios to dominant risk classes (Section 4.5.3), will fail here.

Given the failure of the second approach, we consider the first approach which relies on decomposing the portfolio by shocking the equity and fx risk factors separately. We run 10k scenarios on the stock value keeping the fx rate constant and repeat the exercise for the fx rate, keeping the stock value constant. We calculate the ES for each run and feed the results into the aggregation function using correlation=0.7. The results are as follows:

ES(EQ)	ES(FX)	ES(Aggregated)	ES(Integrated)	% Diff
70.60	71.62	131.12	102.15	28.4%

Table 3: Integrated vs aggregated comparison, correlation=0.7

In this example, the aggregation does a poor job of measuring the compounding effect in the portfolio and, in fact, overestimates the true effect by close to 30%. Switching the correlation to -0.7 should produce a diversification effect since as the stock price falls (rises), the fx rate will rise (fall) helping to stabilize the value of the portfolio. Re-running the analysis above returns the following results:

ES(EQ)	ES(FX)	ES(Aggregated)	ES(Integrated)	% Diff
70.40	71.16	54.83	61.45	-10.8%

[25] This can be determined by breaking the position into two components, one representing the fx risk and the other representing the equity risk and seeing how each component contributes to the total.

Table 2: Integrated vs aggregated comparison, correlation=-0.7

In this case, the aggregation overstates the diversification effect producing a lower risk measure than the integrated approach.[26]

Re-running the above analysis while varying the correlation and stock volatility produces the following result:

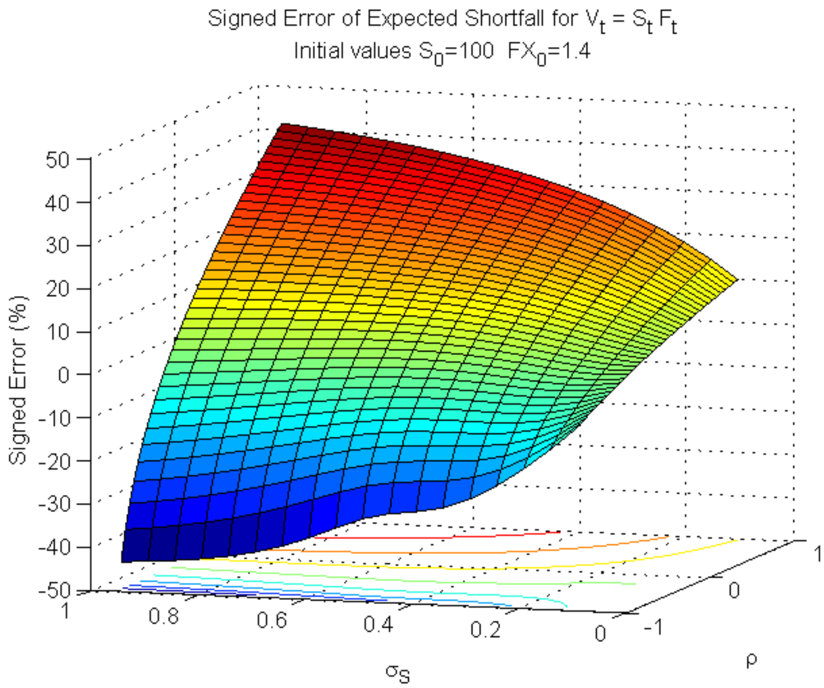

Figure 10: Error comparison between aggregated and integrated model

The most telling feature of the graph in Figure 10 is that the aggregation error significantly grows in areas where the stock volatility is large and the correlation value is extreme (positive or negative). Presumably, this is precisely the case where the Committee intends to use the aggregation (i.e. the stress calibrations will result in large volatilities and the whole point of the exercise is to overlay extreme correlations). **So the Committee has proposed an approach that fails precisely under the conditions where it will be applied.**

In fact, the approach can fail even with moderate adjustments to correlations simply because of the fact that the volatilities themselves may be quite large due to the stressed calibrations. To show this, we re-ran the example above using a correlation of 0.5 but varying the volatilities for both the stock and fx rate:

[26] Note that theoretically, ES(EQ) and ES(FX) should not have changed between both runs. The slight change is a result of sampling error when using 10k scenarios.

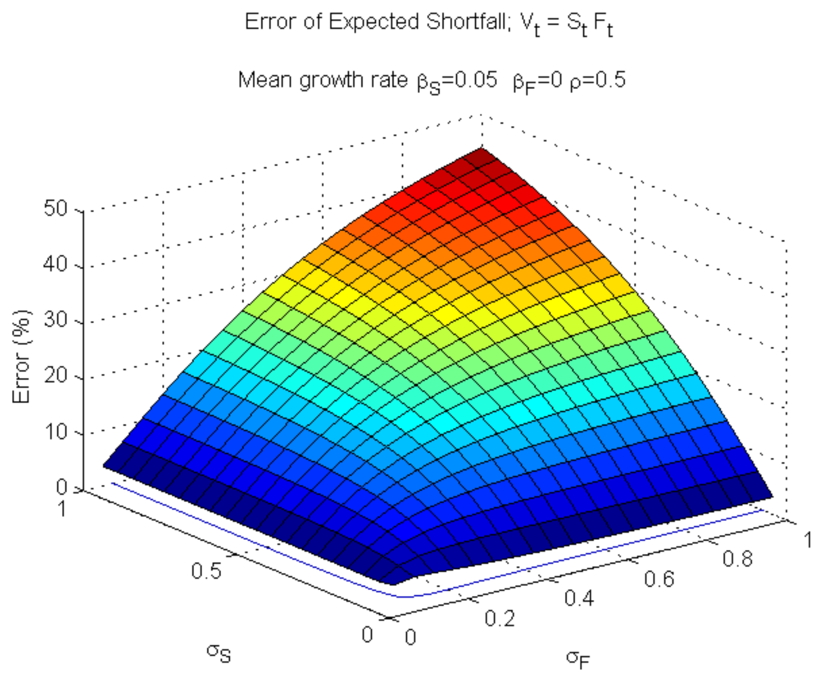

Figure 11: Error comparison between aggregated and integrated model, corr=0.5

Figure11 indicates the aggregation error exceeds 30% even with moderately stressed annual volatilities of 50%. Further examples indicate that the error grows as the time horizon grows. For instance, considering a liquidity horizon of 6m as opposed to 3m produces maximum errors in the range of 80%. **The irony here is that, in an attempt to control diversification benefits, the Committee has elected to use an approach which is known to be unreliable in capturing diversification effects.**

Example 2 - Non-Directional Portfolios Leading to Regulatory Arbitrage

Perhaps the biggest challenge with the Committee's suggested use of the aggregation function has to do with the classification of risk classes. In a word, they are too broad and high-level to be meaningful. For instance, the Committee requires each "L" in the formula to reflect whether a bank is long or short a risk class. Clearly this makes little sense. In a global portfolio, what does being long or short FX actually mean? The bank may have long exposure to some currencies and short exposure to others. As well, given market-making activities for their clients, some banks may even end up with both long and short exposures to the same currency. The same applies to the other risk classifications (Interest Rates, Equities, Commodities, Credit Spreads). Such high-level aggregations will make it virtually arbitrary to determine whether the "L" for each risk class should be +1 or -1. Of course, this problem can be remedied by breaking down the higher level categorizations into more detailed buckets (i.e., US Interest Rates, US Equities etc.) but that leads to many more categories for which the Committee would need to specify inter-group correlations. It also does not completely solve the issue since any level of aggregation continues to make the assignment of directional risks vague (for instance, even with a US Interest Rates

category a bank could be long the short end of the curve and short the long end). This point should not be overlooked. **The ambiguity of whether banks are long or short an aggregate risk class will lead to regulatory arbitrage**. In fact, it would not be difficult to imagine simple optimization routines solving for the combination of "L" values for each risk class that minimizes the aggregation result thereby increasing the likelihood it never comes into play for calculating capital. In this case, regulators would not be able to prove that the allocation of L was an inappropriate reflection of the directional risk. In what follows, we show this arbitrage opportunity using a simple, 2-factor example.

For simplicity, let's assume the bank's total portfolio is broken down into 2 portfolios. Portfolio1 maps to the FX category and its profit and loss profile depends on a single FX factor in the following manner:

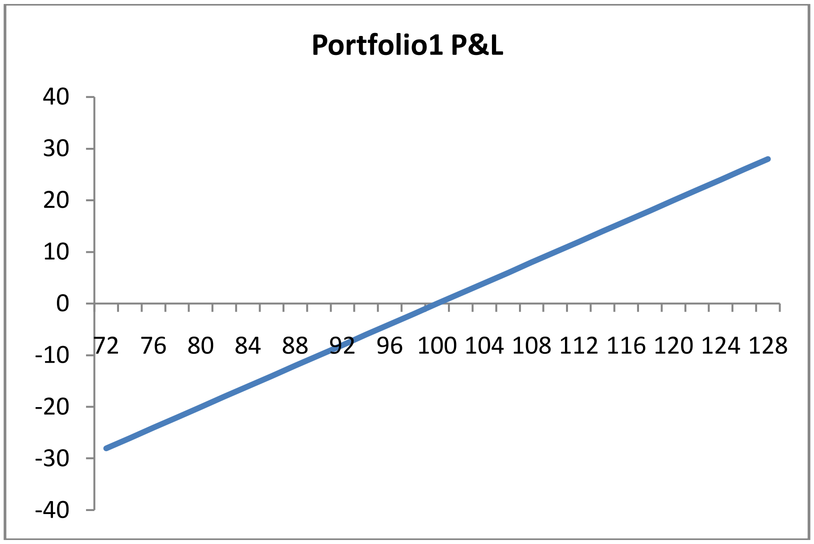

Figure 12: P&L Profile as function of FX moves

Figure 12 clearly indicates that the bank is "long" the FX risk class in this case as increases (decreases) in the risk factor cause portfolio gains (losses).

Portfolio2 maps to the EQ category and its profit and loss profile depends on a single EQ factor in the following manner:

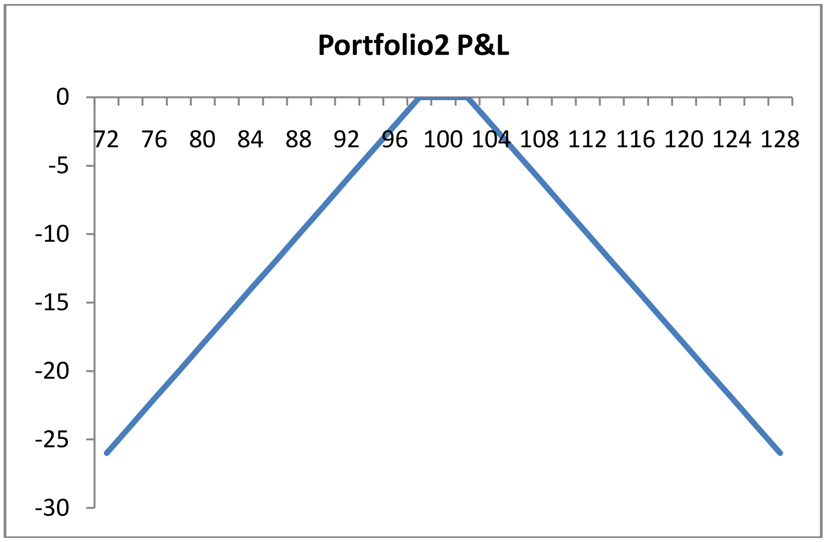

Figure 13: P&L Profile as function of EQ moves

Figure 13 indicates that Portfolio2 is neither long nor short the EQ risk class. In this case, the portfolio's ES measure will receive contributions from both increases and decreases in the EQ risk factor.[27] In this case, allocating a +1 or -1 to the "L" parameter for the ES measure of this portfolio is rather arbitrary. However, an allocation of +1 or -1 will dramatically impact the top-down aggregation.

As before, we assume the EQ and FX factors are simulated using the GBM models in (1) and (2) above with a 3m horizon and the following parameters:

S_0= current stock value = 100
F_0= current fx rate = 100
β_S = drift on the stock = 0%
β_F= drift on the fx rate = 0%
σ_S = annual stock volatility = 23.7%
σ_F = annual fx volatility = 23.7%
W_t^S and W_t^F are standard Brownian motions correlated with ρ = 0.7

Simulating the distribution for Portfolio1 and Portfolio2 independently, we calculate their individual ES and plug into the aggregation assuming an "L" for Portfolio2 of +1. We also simulate the two portfolios in an integrated fashion and calculate the integrated ES. The results are as follows:

ES(Portfolio1)	ES(Portfolio2)	ES(Aggregated)	ES(Integrated)	% Diff
28.72	35.11	58.89	51.72	13.9%

[27] Note that while this P&L profile may look exaggerated, it is easily attainable by selling both a call and put on the same stock. In reality, when dealing with many risk factors in a single risk class, moving all the risk factors up and down in a coordinated fashion can produce far stranger (and ambiguously long vs. short) profiles on realistic bank portfolios.

Table 3: Integrated vs aggregated comparison, L(Portfolio2) = +1

Assigning a value of +1 for the "L" for Portfolio2 produces an aggregation result that is 14% higher than the integrated model. However, as discussed, there is no clear reason why Portfolio2 should be considered a long portfolio with respect to the equity risk factor. Recalculating the aggregation this time assuming a value of -1 for the "L" produces the following result:

ES(Portfolio1)	ES(Portfolio2)	ES(Aggregated)	ES(Integrated)	% Diff
28.72	35.11	25.41	51.72	-50.9%

Table 4: Integrated vs aggregated comparison, L(Portfolio2) = -1

In this case, the aggregation comes in 51% lower than the integrated model. Despite the fact that nothing has really changed in terms of the true, integrated risk profile of the portfolio, the risk result produced by the aggregation has been made irrelevant (i.e., no impact on regulatory capital) simply though the arbitrary assignment of "L".

Re-running the example with varying volatilities for both the stock and fx yields the following graph:

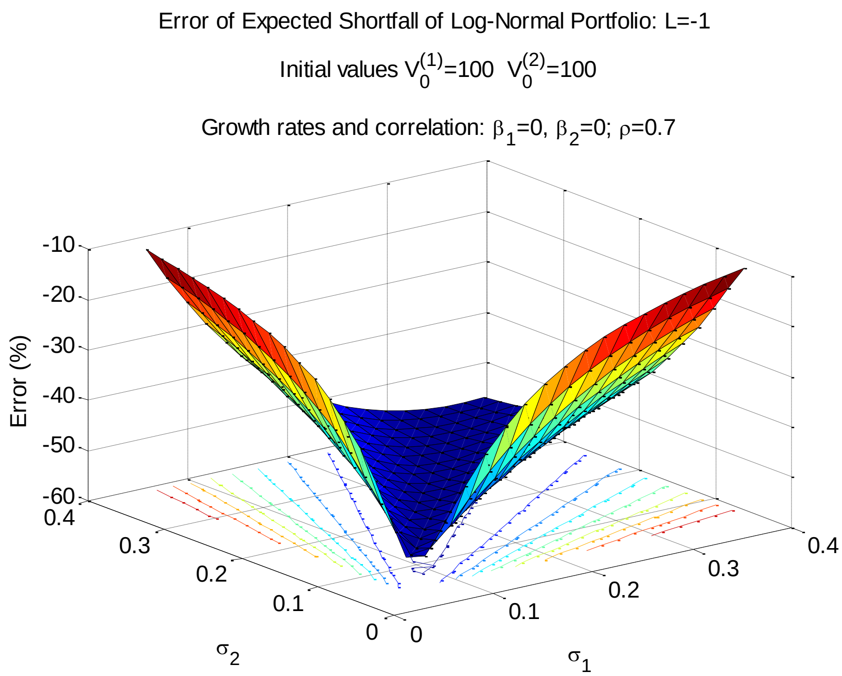

Figure 14: Error comparison between aggregated and integrated model, corr=0.7 and L=-1

Figure 14 has a very interesting shape. As the volatilities become similar (regardless of their value), the two portfolios contribute equally to the integrated ES, causing the aggregation to underestimate the integrated measure by a wide margin. As one risk factor dominates the other (i.e., one volatility is much higher than the other), the aggregation error declines as the joint impacts become less critical. **Extrapolating, we can predict that the more balanced the overall bank portfolio (i.e., the more evenly the individual risk classes contribute to overall**

risk), the greater the opportunity to arbitrage the final risk measure simply by judiciously setting the "L" factors.

Summary

This annex has highlighted numerous concerns and deficiencies associated with the use of top-down aggregations. Given all the shortcomings expressed in the Overview, the use of the top-down aggregation presents real consequences for banks in terms of operational risks and performance challenges. It also presents challenges to the Committee for proper specification of the inter-group correlations to ensure valid matrices. As well, Example 1 indicates clearly that top-down aggregations are not reliable indicators of risk when looking at compounding and diversification effects in the portfolio. It does seem odd that, for the purposes of constraining diversification benefits, the Committee would explicitly adopt a technique that is known to be a failure in modelling diversification effects. Moreover, this is a failure that the Committee itself has acknowledged many years back (see response to question 6 in the letter). Example 2 shows clearly that the Committee's suggestion is rife with regulatory arbitrage as the assignment of "L" for each risk class will be arbitrary given the high-level definition of risk classes.

Overview

The Committee has proposed the use of the top-down aggregation function in formula (1) in an attempt to exert some control over diversification benefits between aggregate risk classes. As stipulated in our response to Question 6 and in Annex B, the use of the function is fraught with challenges and creates opportunities for regulatory capital arbitrage. This Annex discusses an alternative approach which allows the Committee to exert some level of control over the risk factor correlations that banks will use in their internal model. This Annex describes how to "roll up" risk factor correlations into aggregate risk class correlations, how to apply shocks to those aggregate risk class correlations, and then how to cascade those aggregate shocks down to the risk factor level, thereby producing regulatory-adjusted correlation matrices which can be used in the internal model calculations. As stipulated in our response to Question 6, regulatory capital can then be determined by taking the maximum ES calculated using either the original bank correlation matrix or the regulatory-adjusted correlation matrices. Such an approach prohibits the opportunity for any regulatory capital arbitrage.

Background

Let the random variables r_k and r_l denote the returns of risk factors k and l. The correlation of the risk factor returns is

$$\rho_{kl}^{r} = \frac{cov(r_k, r_l)}{\sigma_k \sigma_l}$$

where σ_k and σ_l are the respective return volatilities[28].

Suppose that a bank's portfolio is sensitive to P risk factors. We assume that the internally modeled bank-wide capital charge, $IMCC(C)$, is computed from a joint simulation of all P factors. Thus, it derives from the $P \times P$ variance-covariance (VCV) matrix Ω, where

$$\Omega_{kl} = cov(r_k, r_l) = \rho_{kl}^{r} \sigma_k \sigma_l \,.$$

Ω is calibrated to historical periods of significant financial stress.

Let the portfolio's exposure to risk factor k be m_k. We say that the portfolio is long or short risk factor k if $m_k > 0$ or $m_k < 0$, respectively. Generally, if $m_k m_l \rho_{kl}^{r} < 0$ then the effects of risk factors k and l offset each other to a certain extent. A simulation based on Ω ensures

[28] The superscript "r" on the correlation is used to distinguish this correlation from the correlation in Equation 1 in this annex.

that $IMCC(C)$ reflects the $P(P-1)/2$ pairwise interactions of all P risk factors in times of stress.

Now partition the P risk factors into N groups, or risk classes, where the i^{th} class, C_i, comprises P_i risk factors. Let $IMCC(C_i)$ denote the internally modeled capital charge for risk class C_i. $IMCC(C_i)$ is computed from a simulation based on $\Omega^{(i,i)}$, a $P_i \times P_i$ sub-matrix of Ω that corresponds to the risk factors in C_i. One can attempt to combine the N class-specific capital charges into an overall capital charge in some manner. The regulatory scheme proposes the formula

$$\sqrt{\sum_{i=1}^{N} IMCC^2(C) + \sum_{i=1}^{N} \sum_{j \neq i} \rho_{ij}\ L_i L_j IMCC(C_i) IMCC\left(C_j\right)} \tag{1}$$

where L_i = 1 or –1 if the portfolio is, respectively, long or short risk class C_i, and ρ_{ij} is a prescribed correlation parameter that determines the degree to which the risk associated with risk class C_i offsets that of risk class C_j.

The VCV matrix Ω, which already reflects stressed risk factor correlations, is a more appropriate mechanism for aggregating risks across asset classes in a consistent manner. If the goal of Equation 1 is to further stress the inter-group correlations then we suggest proceeding as follows:

1. Prescribe reasonable shocks to the inter-group correlations, e.g., increase the existing (already-stressed) correlation by some amount so that the matrix remains positive semi-definite (PSD).
2. Implement the shocks by modifying the pairwise risk factor correlations (i.e., the VCV matrix Ω) accordingly.
3. Jointly simulate all P risk factors using the adjusted VCV matrix and compute $IMCC(C)$.

We now outline one possible approach for performing Step 2. The general idea is to adjust the correlations ρ_{kl}^{r}, where risk factors k and l are in different risk classes, so that a benchmark portfolio exhibits the prescribed inter-group correlations. This adjustment method, while conceptually simple, does not guarantee that Ω remains PSD. If the adjusted Ω is not PSD, standard mathematical procedures can be used to minimally modify the pairwise correlations in order to obtain a PSD matrix. Tests with realistic data suggest that even after such corrections, the inter-group correlations are practically equivalent to the prescribed values.

For illustrative purposes, consider Figure 15, which shows a VCV matrix Ω where the P risk factors are partitioned into three risk classes (note that $\Omega^{(i,j)} = \left(\Omega^{(j,i)}\right)^{\mathrm{T}}$ for i = 1, 2, 3 and j = 1, 2, 3). Under the proposed regulatory scheme, $IMCC(C_i)$ is computed from a simulation that is based on $\Omega^{(i,i)}$ for i = 1, 2, 3, and then these class-specific charges are combined using the prescribed correlations ρ_{12}, ρ_{13} and ρ_{23} . Instead, we suggest modifying the off diagonal blocks of Ω (unshaded in Figure 15) so that the inter-group correlations effectively correspond to those prescribed, and then performing a joint simulation of all P risk factors to compute $IMCC(C)$.

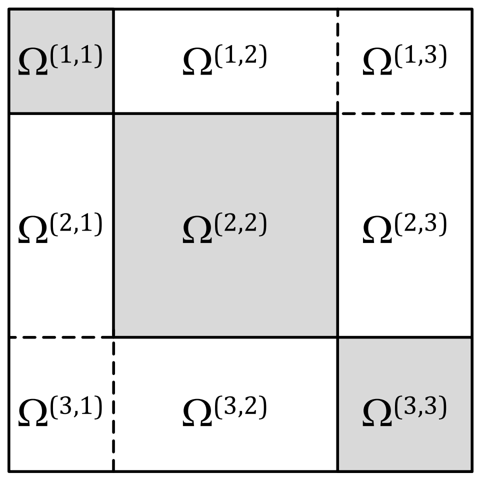

Figure 15: VCV matrix with three risk classes.

Rolling Up Risk Factor Correlations to Group Levels

As noted previously, a correlation between risk classes is meaningful only in a portfolio context; i.e., one cannot claim that "risk classes i and j have correlation ρ_{ij} ", only that "portfolios A and B, with given exposures to risk classes i and j, have correlation $\rho_{ij}(A,B)$". [29] With this in mind, let $\boldsymbol{m}^i$ and $\boldsymbol{m}^j$ denote a set of exposures to risk classes i and j. The correlation between the respective changes in value,

$$d\left(\boldsymbol{m}^i\right) = \sum_{k \in C_i} m_k^i \, r_k$$

[29] This is readily apparent when one considers the correlation between the returns of stocks and bonds, for example. The correlation between the returns of an individual stock and an individual bond can be specified unambiguously. At the asset-class level, however, the correlation necessarily refers to the returns of given stock and bond indexes (i.e., portfolios comprising individual stocks and bonds, respectively, in fixed proportions). Clearly, changing the index weightings also changes the asset-class correlation.

and

$$d(\boldsymbol{m}^j) = \sum_{l \in C_j} m_l^j \, r_l \,,$$

is

$$\begin{aligned}\rho_{ij}\left(\boldsymbol{m}^i, \boldsymbol{m}^j\right) &= \frac{cov\left(d(\boldsymbol{m}^i), d(\boldsymbol{m}^j)\right)}{\sigma(\boldsymbol{m}^i)\sigma(\boldsymbol{m}^j)} \\ &= \frac{\sum_{k \in C_i}\sum_{l \in C_j} m_k^i m_l^j cov(r_k, r_l)}{\sqrt{\sum_{k \in C_i}\sum_{l \in C_i} m_k^i m_l^i cov(r_k, r_l)}\sqrt{\sum_{k \in C_j}\sum_{l \in C_j} m_k^j m_l^j cov(r_k, r_l)}} \,.\end{aligned} \tag{2}$$

Lacking any particular portfolio information, it is reasonable to select a benchmark portfolio having equal exposures to all risk factors in each risk class, i.e., $m_k^i = 1$ for all $k \in C_i$ and $m_k^j = 1$ for all $l \in C_j$. For convenience, write this in vector form as $\boldsymbol{m}^i = \mathbf{1}_{(P_i)}$ and $\boldsymbol{m}^j = \mathbf{1}_{(P_j)}$, where $\mathbf{1}_{(K)}$ denotes a vector of K ones. Then Equation 2 becomes

$$\rho_{ij}\left(\mathbf{1}_{(P_i)}, \mathbf{1}_{(P_j)}\right) = \frac{\sum_{k \in C_i}\sum_{l \in C_j} \rho_{kl}^r \sigma_k \sigma_l}{\sigma_i\left(\mathbf{1}_{(P_i)}\right)\sigma_j\left(\mathbf{1}_{(P_j)}\right)} \,. \tag{3}$$

Cascading Group Adjustments Down to Risk Factor Levels

Suppose now that $\rho_{ij} = \hat{\rho}_{ij}$ is the prescribed correlation of risk classes i and j (hereafter, it is understood that this refers to the correlation exhibited by the equal-exposure benchmark portfolio; to improve readability, we now drop the explicit references to the portfolio). Our goal is to adjust the elements of $\Omega^{(i,j)}$ so that the correlation for risk classes i and j equals $\hat{\rho}_{ij}$. In fact, many different sets of risk factor correlations produce this result. We consider a particularly simple adjustment, which shifts the inherent correlations of all elements of $\Omega^{(i,j)}$ by the same amount, x_{ij}; i.e., for $k \in C_i$ and $l \in C_j$, ρ_{kl}^r is adjusted to

$$\hat{\rho}_{kl}^r = \rho_{kl}^r + x_{ij} \,. \tag{4}$$

From Equation 3 we have

$$\hat{\rho}_{ij} = \frac{\sum_{k \in C_i}\sum_{l \in C_j}(\rho_{kl}^r + x_{ij})\sigma_k \sigma_l}{\sigma_i \sigma_j} = \rho_{ij} + x_{ij}\frac{\sum_{k \in C_i}\sum_{l \in C_j}\sigma_k \sigma_l}{\sigma_i \sigma_j} \tag{5}$$

so that

$$x_{ij} = \frac{(\hat{\rho}_{ij} - \rho_{ij})\sigma_i\sigma_j}{\sum_{k\in C_i}\sum_{l\in C_j}\sigma_k\sigma_l} \tag{6}$$

Thus, the matrix $\widehat{\Omega}$, which exhibits the prescribed correlation $\hat{\rho}_{ij}$ for risk classes *i* and *j*, has elements

$$\widehat{\Omega}_{kl} = \Omega_{kl} + \begin{cases} 0 & if\ k \in C_i, l \in C_j\ and\ i = j \\ \dfrac{(\hat{\rho}_{ij} - \rho_{ij})\sigma_i\sigma_j\sigma_k\sigma_l}{\sum_{k'\in C_i}\sum_{l'\in C_j}\sigma_{k'}\sigma_{l'}} & if\ k \in C_i, l \in C_j\ and\ i \neq j \end{cases}. \tag{7}$$

There is no guarantee that $\widehat{\Omega}$ is actually a valid VCV matrix since it may not be PSD. In this case, we apply a standard algorithm (Higham 2002) that modifies the elements of $\widehat{\Omega}$ in order to produce the closest (in some norm) new matrix $\widetilde{\Omega}$ that is PSD. We choose an equal weighting (corresponding to a Euclidean norm) in the corrective algorithm. As demonstrated in the following experiments, the procedure works well in practice.

The Data

We consider a set of *P* = 759 risk factors that are partitioned into five classes:

- commodities (CO) - 12 factors
  - *Energy* (6): WTI Crude, Brent Crude, Natural Gas, Heat Oil, NY RBOB, Gas Oil
  - *Agriculture* (2): Wheat, Corn
  - *Metals* (4): Gold, Silver, Platinum, Palladium

- credit spreads (CS) - 82 curves × 6 nodes (1y, 3y, 5y, 7y, 10y, 20y) = 492 factors
  - *Banks* (36): Bank Of America, Banca Intesa Sanpaolo, Banca Popolare Milano, BancaUBI, Banco Bilbao Vizcaya Argentaria, Banco Popolare, Banque PSA, Barclays Bank, Bayrische Landesbank, BNP Paribas, Citigroup, Commerzbank, Credit Agricole, Credit Suisse, Danske Bank, Deutsche Bank, Dexia Crediop, DnB NOR Bank, Goldman Sachs, HSBC Bank, ING Bank, JP Morgan Chase, Landesbank Baden Wuerttemberg, Mediobanca, Monte Paschi Siena, Morgan Stanley, Nomura, Nordea Bank, Rabobank, Royal Bank Of Scotland, Santander, SNS Bank, Societe Generale, Svenska Handelsbanken, UBS, Unicredit
  - *Corporates* (28): Allianz, Ally Financial, Assicurazioni Generali, BMW, Compagnie De Saint Gobain, Daimler, Deutsche Telekom, E.ON, Electricite De France, ENEL, ENI, Fiat, Finmeccanica, Ford, France Telecom, GDF Suez,

General Electric, IBM, Koninklijke KPN, Pirelli, Renault, RWE, Telecom Italia, Telefonica, Total, Toyota, Vodafone, Volkswagen
    - *Sovereigns* (18): Austria, Belgium, Denmark, Finland, France, Germany, Ireland, Italy, Japan, Netherlands, Norway, Portugal, Slovenia, Spain, Sweden, Switzerland, United Kingdom, USA

- equity market indices (EQ) - 55 factors
    - Austria (ATS), United Arab Emirates (AED), Argentina (ARS), Australia (AUD), Belgium (BEF), Bulgaria (BGN), Brazil (BRL), Canada (CAD), Switzerland (CHF), Chile (CLP), China (CNY), Colombia (COP), Czech Republic (CZK), Germany (DEM), Denmark (DKK), Egypt (EGP), Spain (ESP), Finland (FIM), France (FRF), United Kingdom (GBP), Greece (GRD), Hong Kong (HKD), Hungary (HUF), Indonesia (IDR), Ireland (IEP), Israel (ILS), India (INR), Italy (ITL), Japan (JPY), South Korea (KRW), Kuwait (KWD), Kazakhstan (KZT), Lithuania (LTL), Morocco (MAD), Mexico (MXN), Malaysia (MYR), Netherlands (NLG), Norway (NOK), New Zealand (NZD), Oman (OMR), Peru (PEN), Philippines (PHP), Pakistan (PKR), Poland (PLN), Portugal (PTE), Quatar (QAR), Russia (RUB), Sweden (SEK), Singapore (SGD), Thailand (THB), Turkey (TRY), Taiwan (TWD), United States (USD), Vietnam (VND), South Africa (ZAR)

- foreign exchange (FX) - 56 factors
    - Australia (AUD), United Arab Emirates (AED), Netherlands Antilles (ANG), Argentina (ARS), Bulgaria, (BGN), Bahrain (BHD), Brazil (BRL), Canada (CAD), Switzerland (CHF), Chile (CLP), China (CNY), Colombia (COP), Czech Republic (CZK), Denmark (DKK), Ecuador (ECS), Egypt (EGP), Euro Area (EUR), United Kingdom (GBP), Hong Kong (HKD), Croatia (HRK), Hungary (HUF), Indonesia (IDR), Israel (ILS), India (INR), Iceland (ISK), Japan (JPY), South Korea (KRW), Kuwait (KWD), Cayman Islands (KYD), Kazakhstan (KZT), Lithuania (LTL), Latvia (LVL), Morocco (MAD), Mexico (MXN), Malaysia (MYR), Norway (NOK), New Zealand (NZD), Oman (OMR), Peru (PEN), Philippines (PHP), Pakistan (PKR), Poland (PLN), Quatar (QAR), Romania (RON), Russia (RUB), Saudi Arabia (SAR), Sweden (SEK), Singapore (SGD), Thailand (THB), Turkey (TRY), Taiwan (TWD), Ukraine (UAH), Uruguay (UYU), Venezuela (VEF), Vietnam (VND), South Africa (ZAR)

- interbank curves (IR) - 24 curves × 6 nodes (3m, 6m, 1y, 2y, 5y, 10y) = 144 factors
    - Australia (AUD), Bulgaria (BGN), Brazil (BRL), Canada (CAD), Switzerland (CHF), Czech Republic (CZK), Denmark (DKK), Euro Area (EUR), United Kingdom (GBP), Hong Kong (HKD), Hungary (HUF), Japan (JPY), South Korea (KRW), Mexico (MXN), Norway (NOK), New Zealand (NZD), Philippines(PHP), Poland (PLN), Russia (RUB), Sweden (SEK), Singapore (SGD), United States (USD), Vietnam (VND), South Africa (ZAR)

The 759 × 759 VCV matrix, Ω, for these risk factors is estimated from their daily log returns over a one-year period. Our experiments consider three such periods, namely 2009, 2010 and 2011. For each year, we compute the 10 inter-group correlations from Ω using Equation 3.

		CS	EQ	FX	IR
2009	CO	-0.3248	0.5694	0.5688	0.1376
	CS		-0.6563	-0.3714	-0.1113
	EQ			0.6070	0.1836
	FX				0.0319
2010	CO	-0.4076	0.6263	0.3796	0.1207
	CS		-0.7415	-0.4646	-0.1483
	EQ			0.5540	0.2455
	FX				0.0593
2011	CO	-0.3578	0.5874	0.3015	0.1436
	CS		-0.6903	-0.4906	-0.2314
	EQ			0.4848	0.2513
	FX				0.1057

Table 4: Initial inter-group correlations

Table 4 shows the initial inter-group correlations for each of the three years considered. Based on the equal-exposure benchmark portfolio, the most significant correlation (approximately -0.70) is between credit spreads and equities.

Correlation Adjustment Procedure

For each year, we now do the following:

1. Given a set of shocks on Ω, obtain the elements of the adjusted risk factor VCV matrix $\widehat{\Omega}$ using Equation 7.
2. Adjust $\widehat{\Omega}$ as necessary to obtain a matrix $\widetilde{\Omega}$ that is PSD.
3. Compute the 10 inter-group correlations from $\widetilde{\Omega}$ using Equation 3 and compare them to the target correlations (taken from Ω adjusted for the shocks) .

This procedure allows us to determine the impact of the PSD adjustment as well as the reasonableness of the cascading (Equation 7) and roll-up (Equation 3) algorithms. The goal is to ensure that once the target correlations are cascaded, adjusted for PSD, and re-aggregated, the resulting aggregate matrix is still very close to the target matrix. In assessing the accuracy (i.e., with respect to meeting the targets) of the procedure, two points must be kept in mind: (i) for very small (in absolute value) correlations, percentage error can be misleading and the actual size of the error should be considered instead; (ii) relative to the target correlation, achieving a more extreme result is not as serious as achieving a less extreme result; thus the signs of the error and the shock, should be considered together. In what follows, we consider two types of target matrices: one based

on specifying arbitrary correlation levels and the other based on applying relative shocks to the original matrix.

Specifying Arbitrary Target Levels

Section 4.5.6 in the Review suggests that the Committee may attempt to constrain diversification through the prescription of supervisory-specified values for each of the 10 inter-group correlations. Setting the target correlations, in Equation 2, to some arbitrary value(s), with complete disregard to the structure (e.g., the signs of the entries of) the original, inter-group matrix is rather naive and inappropriate as it can imply unrealistic risk factor relationships. To emphasize this point, we consider the 2011 matrix and apply the adjustment algorithm assuming supervisory-specified values of either 0.6 or 0.8 across all 10 inter-group correlations. Revisiting the original matrix in Table 1 indicates the inappropriateness of such an approach. Not only do values of 0.6 or 0.8 dramatically increase the actual inter-group correlations experienced during the period (in some cases many multiples of an increase in magnitude) but they also would cause large swings in the sign of several inter-group correlations (for instance, the -0.69 correlation between EQ and CS now becomes 0.8). Table 2 indicates the result of taking the 0.6 or 0.8 correlation levels, cascading them down to risk factor levels using Equation 7, adjusting them for PSD, and then rolling up the resultant correlations using Equation 3.

2011		0.60				0.80			
		CS	EQ	FX	IR	CS	EQ	FX	IR
Target	CO	0.6000	0.6000	0.6000	0.6000	0.8000	0.8000	0.8000	0.8000
	CS		0.6000	0.6000	0.6000		0.8000	0.8000	0.8000
	EQ			0.6000	0.6000			0.8000	0.8000
	FX				0.6000				0.8000
Result	CO	0.4447	0.6915	0.6820	0.5694	0.5717	0.8200	0.8096	0.6925
	CS		0.3981	0.4465	0.4712		0.5119	0.5624	0.5827
	EQ			0.6822	0.5955			0.8191	0.7281
	FX				0.5515				0.6503
Difference	CO	-0.1553	0.0915	0.0820	-0.0306	-0.2283	0.0200	0.0096	-0.1075
	CS		-0.2019	-0.1535	-0.1288		-0.2881	-0.2376	-0.2173
	EQ			0.0822	-0.0045			0.0191	-0.0719
	FX				-0.0485				-0.1497
% Difference	CO	-25.9%	15.3%	13.7%	-5.1%	-28.5%	2.5%	1.2%	-13.4%
	CS		-33.6%	-25.6%	-21.5%		-36.0%	-29.7%	-27.2%
	EQ			13.7%	-0.8%			2.4%	-9.0%
	FX				-8.1%				-18.7%

Table 5: Results for 2011 assuming specified inter-group correlations levels of either 0.6 or 0.8

The large % differences between the final rolled up matrix and the target matrix indicate that the overlay of 0.6 or 0.8 correlations to the cascading function requires such a

significant departure from the original risk factor correlation entries that the matrix is far from PSD. Producing a valid matrix can only be achieved with a significant PSD adjustment which, when rolled back up to an aggregate level, produces significant differences from the regulatory-specified targets. This result implies that the regulatory-specified views (in this example) are incompatible with the real-world data. While the Committee may be indifferent to such a finding since the proposal resorts to using top-down aggregations that are unrelated to the underlying risk class distributions, it does bring into question the Committee's proposal because regulatory capital will be based on a scheme that does not represent the world that existed or could have existed (even in the most stressed of times).

Specifying Relative Shocks

Given the issues associated with specifying arbitrary inter-group correlation levels, it would perhaps be more appropriate to apply relative shocks to the original matrix. This would ensure that the final adjusted matrix retains a semblance of its original structure. For instance, rather than specifying arbitrary correlations, the Committee may consider increasing the inter-group correlations by 20% over their current levels. Since there is no prior knowledge of portfolio structure, the Committee may also consider decreasing the inter-group correlations by 20% and then choosing the maximum ES produced by either matrix. The tables below indicate the results of running the Correlation Adjustment Procedure when applying +20% and -20% shocks to the matrices from 2009, 2010, and 2011.

2009		20% Up				20% Down			
		CS	EQ	FX	IR	CS	EQ	FX	IR
Target	CO	-0.2598	0.6832	0.6826	0.1651	-0.3898	0.4555	0.4551	0.1101
	CS		-0.5251	-0.2971	-0.0890		-0.7876	-0.4457	-0.1335
	EQ			0.7284	0.2203			0.4856	0.1469
	FX				0.0383				0.0255
Result	CO	-0.2563	0.6697	0.6668	0.1605	-0.3804	0.4691	0.4627	0.1101
	CS		-0.5229	-0.2953	-0.0892		-0.7715	-0.4401	-0.1332
	EQ			0.7102	0.2167			0.4921	0.1461
	FX				0.0425				0.0259
Difference	CO	0.0035	-0.0135	-0.0158	-0.0046	0.0094	0.0136	0.0076	0.0000
	CS		0.0022	0.0018	-0.0002		0.0161	0.0056	0.0003
	EQ			-0.0182	-0.0037			0.0065	-0.0007
	FX				0.0042				0.0004
% Difference	CO	-1.4%	-2.0%	-2.3%	-2.8%	-2.4%	3.0%	1.7%	0.0%
	CS		-0.4%	-0.6%	0.2%		-2.0%	-1.3%	-0.2%
	EQ			-2.5%	-1.7%			1.3%	-0.5%
	FX				11.0%				1.7%

Table 6: Results for 2009

2010		20% Up				20% Down			
		CS	EQ	FX	IR	CS	EQ	FX	IR
Target	CO	-0.3261	0.7516	0.4555	0.1448	-0.4891	0.5011	0.3037	0.0965
	CS		-0.5932	-0.3717	-0.1187		-0.8898	-0.5575	-0.1780
	EQ			0.6648	0.2946			0.4432	0.1964
	FX				0.0712				0.0475
Result	CO	-0.3222	0.7223	0.4552	0.1439	-0.4765	0.5205	0.3147	0.1035
	CS		-0.5894	-0.3710	-0.1191		-0.8542	-0.5506	-0.1754
	EQ			0.6530	0.2944			0.4519	0.2026
	FX				0.0720				0.0518
Difference	CO	0.0039	-0.0293	-0.0004	-0.0009	0.0126	0.0194	0.0110	0.0069
	CS		0.0038	0.0006	-0.0005		0.0356	0.0069	0.0026
	EQ			-0.0118	-0.0003			0.0087	0.0062
	FX				0.0008				0.0044
% Difference	CO	-1.2%	-3.9%	-0.1%	-0.6%	-2.6%	3.9%	3.6%	7.2%
	CS		-0.6%	-0.2%	0.4%		-4.0%	-1.2%	-1.5%
	EQ			-1.8%	-0.1%			2.0%	3.2%
	FX				1.2%				9.2%

Table 7: Results for 2010

2011		20% Up				20% Down			
		CS	EQ	FX	IR	CS	EQ	FX	IR
Target	CO	-0.2863	0.7049	0.3618	0.1723	-0.4294	0.4699	0.2412	0.1149
	CS		-0.5522	-0.3925	-0.1851		-0.8283	-0.5888	-0.2777
	EQ			0.5817	0.3015			0.3878	0.2010
	FX				0.1269				0.0846
Result	CO	-0.2843	0.6896	0.3621	0.1731	-0.4215	0.4820	0.2478	0.1266
	CS		-0.5508	-0.3923	-0.1853		-0.8027	-0.5837	-0.2712
	EQ			0.5773	0.2992			0.3936	0.2179
	FX				0.1280				0.0917
Difference	CO	0.0020	-0.0153	0.0003	0.0008	0.0079	0.0121	0.0066	0.0118
	CS		0.0015	0.0002	-0.0002		0.0256	0.0051	0.0065
	EQ			-0.0044	-0.0023			0.0058	0.0169
	FX				0.0012				0.0072
% Difference	CO	-0.7%	-2.2%	0.1%	0.5%	-1.9%	2.6%	2.7%	10.3%
	CS		-0.3%	0.0%	0.1%		-3.1%	-0.9%	-2.3%
	EQ			-0.8%	-0.8%			1.5%	8.4%
	FX				0.9%				8.5%

Table 8: Results for 2011

For all three years (Table 6 - Table 8) the errors are under 5% (in absolute value), except for a few cases in which the original and target correlations are very small in absolute value. The small errors indicate that specifying relative shocks to the already-stressed

correlation matrix may be the more suitable alternative for the Committee to consider when trying to exert some meaningful control over diversification benefits.

Summary

This Annex described a simple procedure for: i) rolling up risk factor correlations into aggregate risk class correlations, ii) applying relative shocks to those aggregate risk class correlations, and iii) cascading those aggregate shocks down to the risk factor level. This procedure produces regulatory-adjusted correlation matrices which can be used directly in the internal model calculations and acts as a more suitable way for the Committee to exert some control over diversification benefits than through an aggregation-based scheme.

We emphasize that the procedure above represents only one possible approach for adjusting the risk factor VCV matrix, Ω, to exhibit the prescribed correlations at the risk-class level. In particular, we have assumed that:

1. A benchmark portfolio having equal exposures to all risk factors in a given class is appropriate for determining the risk class correlations.
2. All risk factor correlations in $\Omega^{(i,j)}$ are adjusted by (but not to) the same amount.

Clearly, one or both of these assumptions can be modified. For example, one could consider replacing the benchmark portfolio with the bank's actual portfolio. A more sophisticated adjustment procedure, such as one that always produces a PSD matrix, is also an interesting possibility for future research.

References

Higham, N.J. (2002), "Computing the Nearest Correlation Matrix - A Problem from Finance," *IMA Journal of Numerical Analysis* 22.

Overview

The Committee proposes to replace VaR with ES in order to better capture tail risk. There is a perception that ES is just as stable as VaR (except for certain, contrived cases, such as extremely fat-tailed loss distributions). In other words, for practical purposes, when estimating risk from simulated or historical loss samples, ES produces a similar sampling error to VaR. In fact, this is incorrect. As shown below, ES estimates have a larger standard error than VaR estimates when, for instance, losses are Normally distributed. Thus, when selecting an appropriate risk measure for determining regulatory capital requirements, the Committee should consider the fact that ES may often be harder to estimate than VaR.

Standard Error

Let the random variable X, with distribution function F, represent the monetary loss of some portfolio. The portfolio's risk is measured by the VaR and the ES at quantile level α, denoted respectively by VaR_α and ES_α, where $\alpha > 0.5$ (i.e., we are interested in the right tail of the distribution). Suppose that the risk is estimated from a sample $\boldsymbol{x}$ of n losses, drawn at random from F, and let $x^{(k)}$ denote the kth order statistic in the sample (i.e., $x^{(1)} \leq x^{(2)} \leq \ldots \leq x^{(n)}$). To simplify the discussion, assume that n is sufficiently large that the product $m = n\alpha$ is an integer.

A standard estimate of VaR_α is the sample quantile

$$\widehat{VaR}_\alpha = x^{(m)}\,, \tag{8}$$

while ES_α is estimated as the average of the losses exceeding the sample quantile

$$\widehat{ES}_\alpha = \frac{1}{n(1-\alpha)} \sum_{k=m+1}^{n} x^{(k)}\,. \tag{9}$$

The asymptotic ($n \to \infty$) variances of these estimates are (see, for example, Kerkhof et al 2002, Manistre and Hancock 2005)

$$var\left(\widehat{VaR}_\alpha\right) = \frac{\alpha(1-\alpha)}{n[f(VaR_\alpha)]^2}\,, \tag{10}$$

where f is the probability density function, and

$$var\left(\widehat{ES}_\alpha\right) = \frac{var(X|X > VaR_\alpha) + \alpha(ES_\alpha - VaR_\alpha)^2}{n(1-\alpha)}\,. \tag{11}$$

Suppose that losses are normally distributed, i.e., $X \sim \mathrm{N}(\mu, \sigma)$. Let Φ and ϕ denote the distribution and density functions, respectively, of the standard normal distribution. In this case,

$$VaR_\alpha = \mu + Z_\alpha \sigma \,, \tag{12}$$

where

$$Z_\alpha = \Phi^{-1}(\alpha) \tag{13}$$

and

$$ES_\alpha = \mu + K_\alpha \sigma \,, \tag{14}$$

where

$$K_\alpha = \frac{1}{1-\alpha} \int_{Z_\alpha}^{\infty} z\phi(z)dz = \frac{\phi(Z_\alpha)}{1-\alpha}. \tag{15}$$

Since

$$f(VaR_\alpha) = \frac{\phi(Z_\alpha)}{\sigma} = \frac{K_\alpha(1-\alpha)}{\sigma}, \tag{16}$$

it follows that

$$var\left(\widehat{VaR}_\alpha\right) = \frac{\alpha\sigma^2}{nK_\alpha^2(1-\alpha)}. \tag{17}$$

It can be shown that

$$var\left(\widehat{ES}_\alpha\right) = \sigma^2 \left[\frac{Z_\alpha K_\alpha + 1 - K_\alpha^2 + \alpha(K_\alpha - Z_\alpha)^2}{n(1-\alpha)}\right]. \tag{18}$$

Thus, the ratio of the standard errors of $\widehat{ES}_\alpha$ and $\widehat{VaR}_\alpha$ is

$$\frac{SE\left(\widehat{ES}_\alpha\right)}{SE\left(\widehat{VaR}_\alpha\right)} = K_\alpha \sqrt{\frac{Z_\alpha K_\alpha + 1 - K_\alpha^2 + \alpha(K_\alpha - Z_\alpha)^2}{\alpha}}. \tag{19}$$

Note that this ratio is the same for any normal distribution, i.e., it is independent of μ and σ. Figure 16 plots this ratio for $0.5 \le \alpha \le 0.999$. It is evident that for $\alpha \ge 0.65$, $\widehat{ES}_\alpha$ has a larger standard error than $\widehat{VaR}_\alpha$.

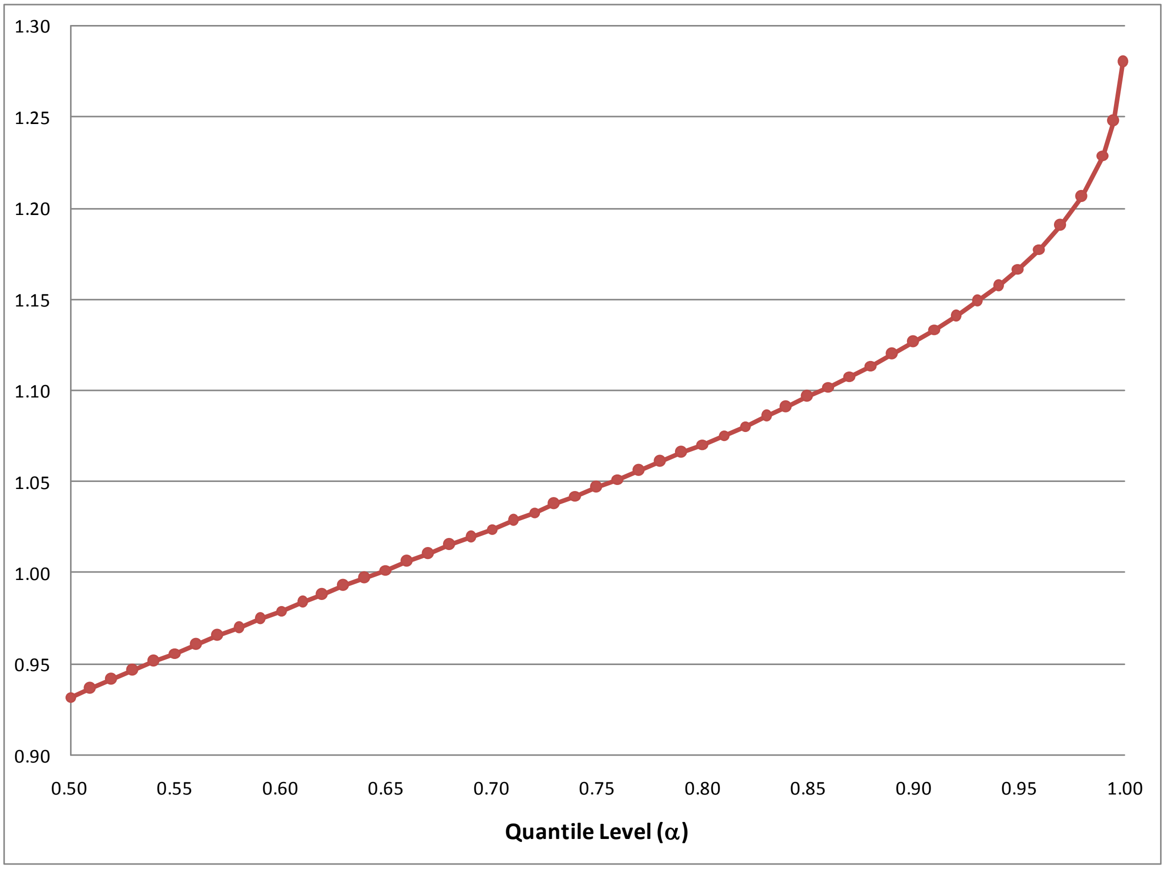

Figure 16: Ratio of standard error of $\widehat{ES}_\alpha$ to that of $\widehat{VaR}_\alpha$ for normal distributions.

When $X \sim N(\mu, \sigma)$, the ratio of the relative standard errors of $\widehat{ES}_\alpha$ and $\widehat{VaR}_\alpha$ is

$$\frac{RSE\left(\widehat{ES}_\alpha\right)}{RSE\left(\widehat{VaR}_\alpha\right)} = \frac{\mu + Z_\alpha \sigma}{\mu + K_\alpha \sigma} \times \frac{SE\left(\widehat{ES}_\alpha\right)}{SE\left(\widehat{VaR}_\alpha\right)}. \quad (20)$$

Unlike Equation 19, Equation 20 depends on μ and σ. Thus, for a given quantile level α, the relative sizes of RSE($\widehat{ES}_\alpha$) and RSE($\widehat{VaR}_\alpha$) vary with the underlying distribution. Note that as μ increases, the ratio of the relative standard errors approaches that of the standard errors.

To validate these results empirically, we estimate $VaR_{0.99}$ and $ES_{0.99}$ for $X \sim N(0, 1)$ using random samples of size n = 1000, 2000, 3000, 4000 and 5000.
Table 7 shows the theoretical (asymptotic) standard errors along with the empirical standard errors as computed from 500 independent trials. The experimental results agree well with the theoretical values, showing the standard error of $\widehat{ES}_{0.99}$ to be approximately

23% larger than that of $\widehat{VaR}_{0.99}$. Similarly, the relative standard error of $\widehat{ES}_{0.99}$ exceeds that of $\widehat{VaR}_{0.99}$ by roughly 7%, as predicted (Table 8).

Table 7: Theoretical and empirical standard error of risk estimates for the standard normal distribution.

Sample Size (n)	$SE(\widehat{VaR}_{0.99})$		$SE(\widehat{ES}_{0.99})$		$SE(\widehat{ES}_{0.99})/SE(\widehat{VaR}_{0.99})$	
	Theoretical	Empirical	Theoretical	Empirical	Theoretical	Empirical
1000	0.1181	0.1181	0.1451	0.1431	1.2291	1.2117
2000	0.0835	0.0841	0.1026	0.0999	1.2291	1.1879
3000	0.0682	0.0642	0.0838	0.0801	1.2291	1.2470
4000	0.0590	0.0574	0.0725	0.0698	1.2291	1.2156
5000	0.0528	0.0505	0.0649	0.0629	1.2291	1.2444

Table 8: Theoretical and empirical relative standard error of risk estimates for the standard normal distribution.

Sample Size (n)	$RSE(\widehat{VaR}_{0.99})$		$RSE(\widehat{ES}_{0.99})$		$RSE(\widehat{ES}_{0.99})/RSE(\widehat{VaR}_{0.99})$	
	Theoretical	Empirical	Theoretical	Empirical	Theoretical	Empirical
1000	0.0507	0.0508	0.0544	0.0537	1.0728	1.0576
2000	0.0359	0.0362	0.0385	0.0375	1.0728	1.0369
3000	0.0293	0.0276	0.0314	0.0301	1.0728	1.0884
4000	0.0254	0.0247	0.0272	0.0262	1.0728	1.0611
5000	0.0227	0.0217	0.0243	0.0236	1.0728	1.0861

When $X \sim$ N(2, 1), the relative standard error of $\widehat{ES}_{0.99}$ is approximately 14% larger than that of $\widehat{VaR}_{0.99}$ (Table 9).

Table 9: Theoretical and empirical relative standard error of risk estimates for the N(2, 1) distribution.

Sample Size (n)	$RSE(\widehat{VaR}_{0.99})$		$RSE(\widehat{ES}_{0.99})$		$RSE(\widehat{ES}_{0.99})/RSE(\widehat{VaR}_{0.99})$	
	Theoretical	Empirical	Theoretical	Empirical	Theoretical	Empirical
1000	0.0273	0.0263	0.0311	0.0292	1.1398	1.1098
2000	0.0193	0.0191	0.0220	0.0214	1.1398	1.1202
3000	0.0158	0.0156	0.0180	0.0184	1.1398	1.1799
4000	0.0136	0.0142	0.0156	0.0164	1.1398	1.1560
5000	0.0122	0.0122	0.0139	0.0145	1.1398	1.1868

In contrast, when $X \sim$ N(-2, 1), $VaR_{0.99}$ = 0.3263 is much closer to zero than $ES_{0.99}$ = 0.6652, which tends to inflate RSE($\widehat{VaR}_{0.99}$). In this case, $\widehat{ES}_{0.99}$ has a lower relative standard error than $\widehat{VaR}_{0.99}$ (Table 10).

Table 10: Theoretical and empirical relative standard error of risk estimates for the N(-2, 1) distribution.

Sample Size (n)	$RSE(\widehat{VaR}_{0.99})$		$RSE(\widehat{ES}_{0.99})$		$RSE(\widehat{ES}_{0.99})/RSE(\widehat{VaR}_{0.99})$	
	Theoretical	Empirical	Theoretical	Empirical	Theoretical	Empirical
1000	0.3617	0.3524	0.2181	0.2059	0.6030	0.5845

2000	0.2558	0.2572	0.1542	0.1493	0.6030	0.5805
3000	0.2089	0.2062	0.1259	0.1250	0.6030	0.6062
4000	0.1809	0.1800	0.1091	0.1096	0.6030	0.6087
5000	0.1618	0.1569	0.0975	0.0990	0.6030	0.6308

Summary

For practical purposes, ES can often be harder to estimate than VaR. This is demonstrated by the fact that, in the case of Normally distributed losses, ES estimates have a larger standard error than VaR estimates. This is contrary to the common perception that the ES measure is just as stable as the VaR measure. For this reason, ES may not be better at capturing tail risk than VaR, assuming the option of using more scenarios is not available.

References

Kerkhof, J., Melenberg, H. and H. Schumacher (2002), “Model Risk and Regulatory Capital,” Discussion paper 2002-27, Department of Economics and Operations Research, Tilburg University.

Manistre, B.J. and G.H. Hancock (2005), “Variance of the CTE Estimator,” North American Actuarial Journal 9:2, 129-156.

Overview

From a backtesting standpoint, we propose to continue using VaR for the purposes of backtesting the frequency of violations. However, to calculate capital, we suggest multiplying VaR by a scaling factor which equals the average ratio (Actual Loss / VaR) across the observations that produce a violation. This is tantamount to assigning capital based on an *ex-post* ES measure. A potential concern is that the ratio will be measured based on a 1d horizon but applied to a much longer horizon, which may not be appropriate. We consider the behaviour of this ratio for three different portfolio processes: normal and lognormal diffusions, and a jump process.[30] In almost all cases, calculating a VaR factor based on 1d horizon should be a conservative measure of the factor for longer horizons.

Normal Process

We assume the dynamics of a portfolio's value through time is described by a normal diffusion process with a drift:

$$S_t = S_0(1 + \beta t + \sigma W_t) \tag{1}$$

where
S_0 = portfolio's value today
S_t = portfolio's value at t
β = drift on portfolio value
σ = annual portfolio volatility
W_t = standard brownian motion

The VaR of this portfolio through time can be expressed as:

$$VaR_p(t) = -\beta S_0 t + \alpha_p \sigma S_0 \sqrt{t} \tag{2}$$

where
α_p = the p-quantile of the standard Normal N(0,1) distribution, $\alpha = \Phi^{-1}(p)$

and the ES can be expressed as:

$$ES_p(t) = -\beta S_0 t + q^{-1}\varphi(\alpha_p)\sigma S_0 \sqrt{t} \tag{3}$$

where
$\varphi(\alpha_p)$ = the standard normal probability density function
q = 1-p

Therefore the ratio:

[30] In what follows, we omit much of the derivation of the formulas and focus primarily on the results and implications. The Committee is welcome to contact us should they choose to see the more detailed derivations.

$$Ratio_p(t) = \frac{ES_p(t)}{VaR_p(t)} \tag{4}$$

can be represented as:

$$Ratio_p(t) = \frac{-\beta S_0 t + q^{-1}\varphi(\alpha_p)\sigma S_0\sqrt{t}}{-\beta S_0 t + \alpha_p \sigma S_0 \sqrt{t}} \tag{5}$$

Notice that the behaviour of the ratio in (5) is fully determined by the coefficient β, representing the growth rate of the portfolio. If it is assumed that β=0, the ratio is constant over any liquidity horizon:

$$Ratio_p(t) = \frac{\varphi(\alpha_p)}{q\alpha_p} \tag{6}$$

If β>0, the factor will grow as the liquidity horizon lengthens. If β<0, the factor calculated at a 1d horizon will be greater than the factor calculated at any longer horizon and hence would be considered a conservative measure of risk.

Log-Normal Process

We assume the dynamics of a portfolio's value through time is described by a log-normal diffusion process with a drift[31]:

$$S_t = S_0 e^{\left(\beta - \frac{\sigma^2}{2}\right)t + \sigma W_t} \tag{7}$$

In this case, the VaR of this portfolio through time can be expressed as:

$$VaR_p(t) = S_0\left(1 - e^{\left(\beta - \frac{\sigma^2}{2}\right)t + \sigma\sqrt{t}\alpha_q}\right) \tag{8}$$

and the ES as:

$$ES_p(t) = S_0\left(1 - \frac{1}{q}e^{\beta t}\Phi(\alpha_q - \sigma\sqrt{t})\right) \tag{9}$$

where Φ = standard normal distribution function

and the ratio as:

[31] See footnote 18 in Annex A.

$$Ratio_p(t) = \frac{1 - \frac{1}{q} e^{\beta t} \Phi(\alpha_q - \sigma\sqrt{t})}{1 - e^{\left(\beta - \frac{\sigma^2}{2}\right)t + \sigma\sqrt{t}\alpha_q}} \tag{10}$$

The ratio is monotonically decreasing for all values of β and σ and for liquidity horizons less than or equal to 1y.[32] In this case, calculating the factor at 1d horizons is a conservative measure of risk.

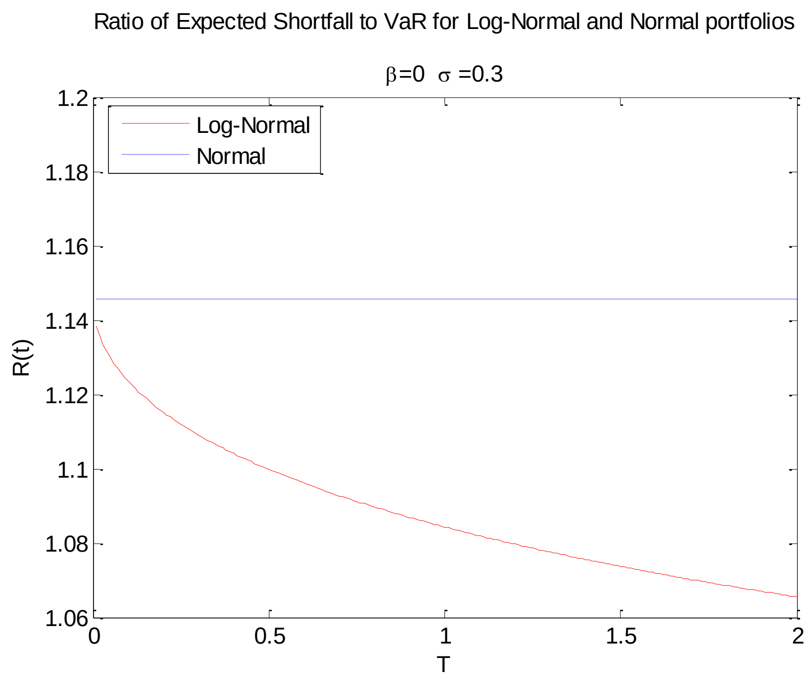

Figure 17: Ratio behaviour for normal and log-normal, β=0

[32] The ratio is not always monotonically decreasing. Usually for σ<0.5 and β>0.75, the ratio will start to rise after t>5years. However, this is irrelevant for the Committee's needs as the longest liquidity horizon is t=1 year.

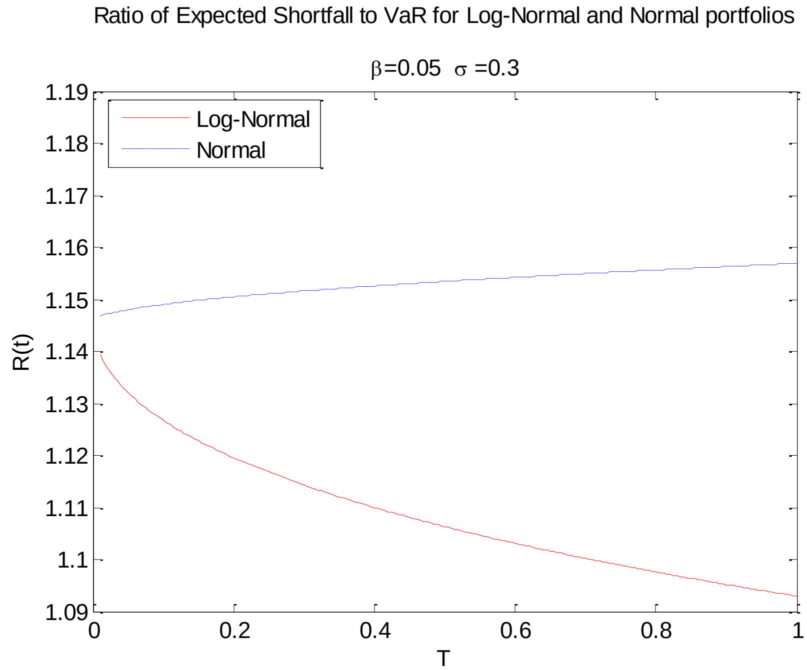

Figure 18: Ratio behaviour for normal and log-normal, β=0.05

Figures 17 and 18 above indicate the behaviour of the ratio for normal and log-normal portfolios with σ=0.3. As discussed, only in the normal case with β>0 will the ratio increase with liquidity horizon.

Varying both the σ and t in the log-normal case produces the following graph:

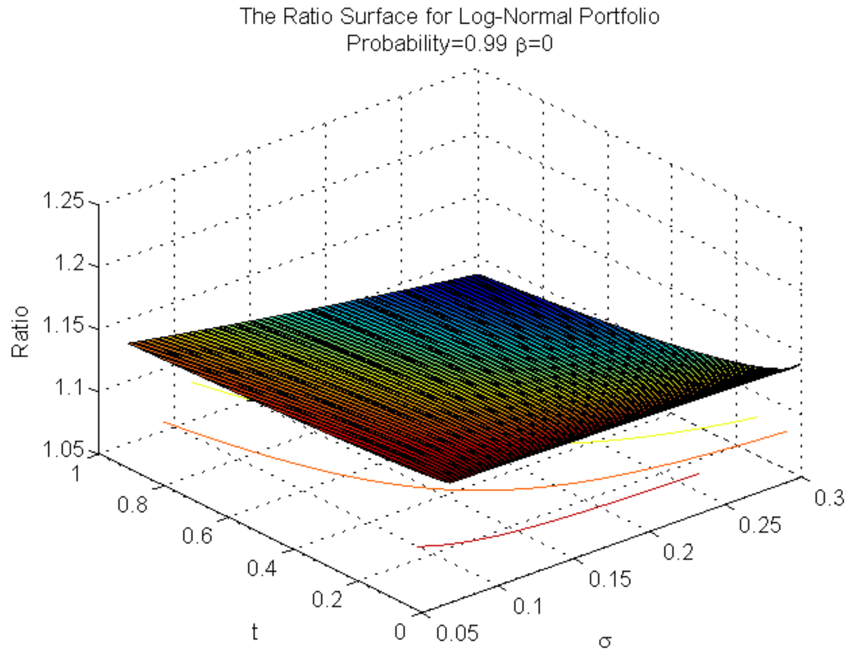

Figure 19: Ratio behaviour for log-normal, varying σ and t

Figure 19 indicates a very flat surface with slight decreasing over the liquidity horizon for varying volatility levels.

Jump Process

We assume the dynamics of a portfolio's value through time is described by a jump model represented as a Compound Poisson Process. The process is described by three parameters: λ representing the intensity of jumps, μ representing the average jump size and σ representing the standard deviation of the jumps. We assume that the jump sizes have a Normal distribution $N(\mu,\sigma^2)$ and that the independent counting process υ_t represents the number of jumps by time t.

Therefore, the cdf of portfolio losses (L) in the Compound Poisson Process above can be represented as:

$$G_t(x) := \mathbb{P}(L_t \leq x) = \sum_{k=0}^{\infty} p_k(t)\Phi\left(\frac{x-k\mu}{\sqrt{k}\sigma}\right) \tag{11}$$

where $p_k(t) := \mathbb{P}(\upsilon_t = k) = e^{-\lambda t}\frac{(\lambda t)^k}{k!}$, $k = 0, 1, 2, \ldots$

Therefore, the $\text{VaR}_p(t)$ of the portfolio can be found by solving the following equation:

$$\sum_{k=0}^{\infty} p_k(t)\Phi\left(\frac{x-k\mu}{\sqrt{k}\sigma}\right)=p \tag{12}$$

and the $ES_p(t)$ can be calculated as follows:

$$ES_p(t)=q^{-1}e^{-\lambda t}\sum_{k=1}^{\infty}\frac{(\lambda t)^k}{k!}\left(\sigma\sqrt{k}\cdot\varphi\left(z_p(k)\right)+k\mu\cdot\Phi\left(z_p(k)\right)\right) \tag{13}$$

where $z_p(k)=\frac{x_p-k\mu}{\sigma\sqrt{k}}$

Based on (12) and (13) above, we can calculate $Ratio_p(t)$ numerically (solving for VaR). The two figures below show the behaviour of the ratio for various time horizons. The first figure varies the intensity of the jump process while the second figure varies the volatility of the jumps themselves.

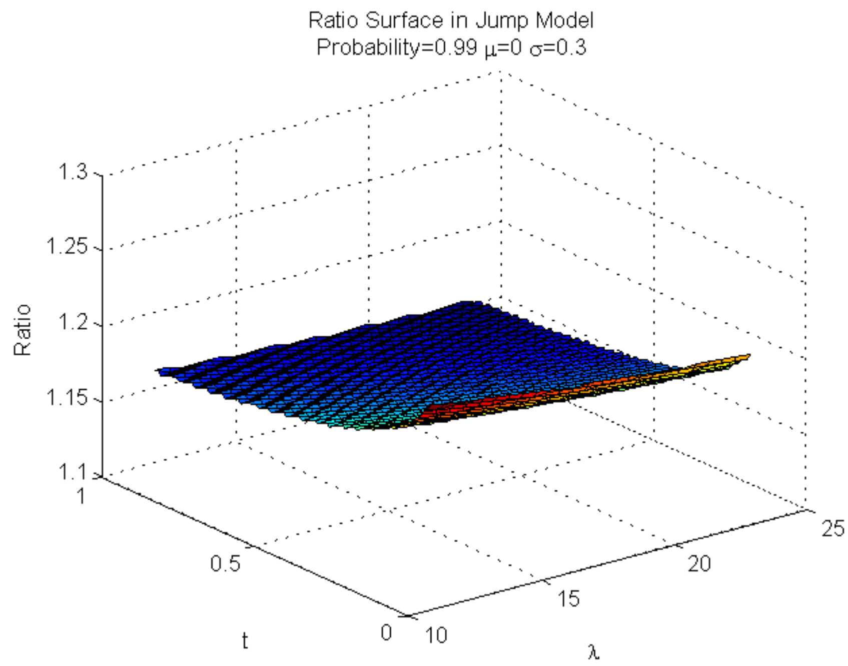

Figure 20: Ratio behaviour for jump process, varying λ

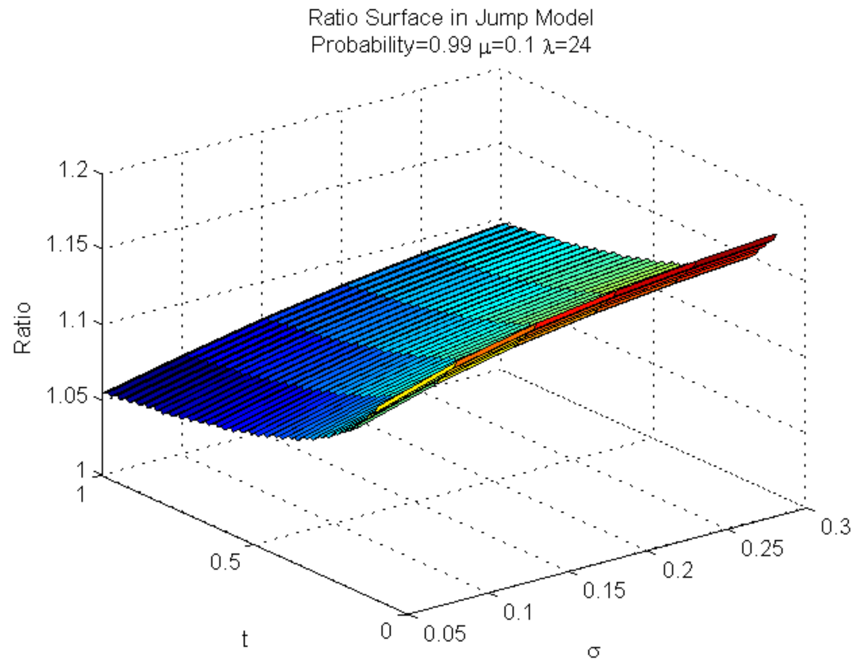

Figure 21: Ratio behaviour for jump process, varying σ

Figures 20 and 21 indicate that the ratio for the Jump process is also relatively flat or slightly decreasing across lengthening liquidity horizons. Again, this implies that calculating a factor based on a 1d horizon and applying it to a longer horizon would be conservative from a regulatory capital standpoint.

Summary

We considered the behaviour of the proposed ratio for three different portfolio processes: normal and lognormal diffusions, and a jump process. Apart from a normal process with positive drift, all other cases indicate that calculating a VaR factor based on 1d horizon should be a conservative measure of the factor for longer horizons.

About Algorithmics, an IBM Company

Algorithmics is a leading provider of risk solutions. Financial organizations from around the world use Algorithmics' software to help them make risk-aware business decisions. Algorithmics' analytics and advisory services assist firms in taking steps towards maximizing shareholder value and meeting regulatory requirements. Supported by a global team of risk experts based in all major financial centers, Algorithmics offers award-winning solutions for market, credit and operational risk, as well as collateral and capital management.
Algorithmics was founded in 1989 and was acquired by IBM on 1st October 2011 to join IBM's Business Analytics division. Algorithmics, has clients in 55 countries, offices in 23 countries, and employs 900 risk professionals. It provides software and analytics solutions in the fields of market, credit, liquidity, and operational risk, and capital and collateral management.

About IBM Business Analytics

IBM Business Analytics software delivers actionable data-driven insights that helps organizations to work smarter and achieve better business outcomes. The comprehensive portfolio offers solutions for business intelligence, predictive analytics, decision management, performance management, risk management and regulatory compliance.
Business Analytics solutions enable companies to identify and visualize trends and patterns in areas that can have a profound impact on business performance. Empowered with analysis, firms can compare scenarios, anticipate potential threats and opportunities, balance risks against expected returns, and better plan, budget and forecast resources. By making analytics widely available, organizations can align tactical and strategic decision-making towards achieving business goals. For further information please visit ibm.com/business-analytics.

