MSCI

7 September 2012

Submitted by email to baselcommittee@bis.org

Secretariat of the Basel Committee on Banking Supervision
Bank for International Settlements
CH-4002 Basel Switzerland

Re: MSCI comments on the Basel Committee for Banking Supervision Consultative Document—Fundamental Review of the Trading Book

To the Committee members:

We are pleased to offer these comments on the Basel Committee's recent Fundamental Review of the Trading Book. We commend the Committee on its thorough review of both industry practices and academic approaches, and for evaluating the experiences of the financial community with risk models during the financial crisis of 2008. We also commend the Committee for its openness and courage in making available for public comment a set of ideas that is not yet a fully formed policy proposal. The industry can only benefit from the dialogue that this document is sure to provoke.

MSCI and RiskMetrics Group have been an active participant in the dialogue between the industry and banking regulators, dating back to the initial discussions in 1994 on what became the 1996 Market Risk Amendment to the Basel Accord. At that time, JP Morgan had just released publically a market risk methodology and dataset—known as RiskMetrics—in part to provide the industry and supervisors with a benchmark for internal market risk methodologies. This release was followed in 1997 by the CreditMetrics methodology to assess banking book capital, which in turn was an integral piece of the dialogue leading to the Basel II agreement. After spinning off from JP Morgan in 1998, RiskMetrics Group continued to develop these methodologies and to provide services to banks and other financial institutions. RiskMetrics Group was acquired by MSCI in 2010, and the firm continues to research and provide these services to the financial community.

While our comments today find areas for specific improvement, we find ourselves in broad agreement with the Committee's high level goals. Our criticism is mostly on implementation rather than motivation. Moreover, we offer our suggestions in the spirit of industry collaboration, and have in all cases tried to recommend alternative approaches where we have disagreed with the Committee's proposals.

7 World Trade Center, 250 Greenwich Street 49th Floor, New York, NY 10007 T: +1.212.804.3900, F: +1.212.804.2919

In the attachment, we provide detailed responses to the Committee's questions, taking the liberty to discuss broader themes than the specific Committee requests. We provide a short summary of our comments below.

- The new, significant role of the standardized model raises questions about the incentives for banks to build and seek approval for internal models. (Question 3)
- The proposed liquidity risk framework is a good start, and the push toward standards on data is helpful. The implementation of the framework, however, is not a realistic treatment of how liquidity issues transpire. We suggest an alternate approach. (Question 2)
- The focus on P&L attribution as a primary step in model validation is an important advance. We suggest ways that this method could be extended beyond the Committee's proposed application. In particular, we recommend that unobserved sources of risk (default, liquidity premia, basis) be identified in this stage and modeled under a common framework. (Question 8)
- The redefinition of capital based on Expected Shortfall (ES) under a stressed calibration is a good choice. While ES itself does not pose significant problems, the lack of clarity on standards for the calibration is a roadblock to robust statistical model validation. (Question 5, Question 8)

In closing, we reiterate our support of the Committee's efforts to introduce a significant step forward in the management of trading book risk, and applaud this insightful document. We appreciate the opportunity to comment, and look forward to continuing to contribute to the dialogue. We are available for further comment or clarification as necessary and welcome future communicaton.

Sincerely,

Christopher C. Finger
Executive Director
Applied Research
christopher.finger@msci.com

Carlo Acerbi
Executive Director
Analytics Research
carlo.acerbi@msci.com

MSCI

Basel Committee on Banking Supervision
Consultative Document: Fundamental Review of the Trading Book
Comments by MSCI Inc.[1]

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

Summary

We agree with the general tenets of the proposed liquidity approach. All things being equal, positions with less liquidity should attract more capital, and the liquidity approach should provide for a smoother transition from trading to banking book treatment of positions. But we disagree with the basic definition of liquidity risk on which the Committee bases its approach, and propose an alternative line of thinking that is more consistent with bank dynamics under liquidity stress. As for the proposal to capitalize for potential shocks to liquidity premia, we see this as unnecessarily specific, and propose that the treatment of liquidity premia be included in a general treatment of (possibly) unmodelable risk factors.

Liquidity parameterization

The proposed approach begins with the requirement that a bank assign each of its holdings to risk factors, and each risk factor to coarse liquidity horizon buckets. The choice of liquidity horizon as the essential parameter is in one sense a reasonable one: other supervisory and disclosure frameworks (for instance, hedge fund disclosure under Form PF in the United States) operate from this same notion. We also support the Committee's proposal to specify floors for the liquidity horizons for certain asset classes. To date, liquidity assessment in many asset classes is largely subjective, and in need of more objective, empirical standards. Any requirement that the specification of liquidity data be empirically based would place significant operational burdens on the banks affected. Thus, even beyond providing floors, the Committee could alleviate a significant hurdle by contributing to an area with few measurement standards. We encourage the Committee to research liquidity in asset classes where there has been little empirical work to date, and where the supervisory community may enjoy privileged access to data.

One criticism we have of the Committee's proposed parameterization is that it is based on risk factors rather than instruments or positions. We feel that this would force banks (as well as the Committee) to make awkward choices about the liquidity profile of factors that exist only in a risk model, but do not represent anything that actually trades. Calibrating the liquidity

[1] For further information, contact Christopher Finger (christopher.finger@msci.com) or Carlo Acerbi (carlo.acerbi@msci.com).

framework will be difficult enough without such contortions. We recommend that the Committee reconsider, and base the liquidity risk framework instead on the assignment of liquidity horizons to classes of instruments or positions.

A second criticism is that the liquidity horizon paradigm by construction neglects trading costs in the sense of bid-ask spreads. The definition of liquidity horizon is the length of time it would take to close (or hedge all risks of) a given position so that trading costs associated with the closeout are immaterial. The bid-ask spread, however, represents a floor to the trading cost which cannot be mitigated by waiting longer to trade or trading in smaller quantities. When the bid-ask spread is itself a material cost—which will be true for many asset classes—the liquidity horizon paradigm ignores it.

Model implementation

The second component of the proposed approach is the implementation of the risk model. Here, the Committee has adopted a *risk-to-liquidation* approach, stipulating that the market risk model include the possible market movements on any given position over the length of time it would take to fully liquidate that position. While we do find this approach conceptually coherent, we have three strong reservations: first, as discussed, the approach neglects the cost associated with bid-ask spreads; second, the approach neglects the effect of position size, as concentration risks are covered under an as yet undefined separate approach for endogenous liquidity effects; and third, the approach does not reflect the reality of the risks that the Committee is seeking to capitalize.

We think it is unrealistic to assume that under stress, a bank would slowly liquidate or hedge its positions, so as to suffer no market impact, and hope that the cash ultimately raised would be sufficient to keep the bank solvent. The reality is that liquidity horizons are imposed from outside, by margin calls, funding shortfalls and other short term demands for cash, as well as from the need to turn over the portfolio to service clients or produce a desired return. As a going concern, a bank must trade, likely suffering from market impact, when faced with these critical demands.

As an alternate to the approach of "stretching" market risk under risk-to-liquidation, we recommend that the Committee consider a separate capital charge for liquidity cost. The basic conceptual framework for this charge would be that a bank's liquidity demands require that it transact some portion of its portfolio in the short term, and that these transactions will come with a cost, that cost being greater for less liquid positions. This cost would be sensitive not only to the exogenous liquidity of the assets in the portfolio, but also to the endogenous effect of large positions, as well as to the need for the portfolio in question to provide liquidity.[2]

[2] A concrete proposal for one such framework is presented in Finger and Acerbi (2010), attached with this submission.

It is true that the liquidity cost framework may be more complex than risk-to-liquidation, but we believe that a realistic, coherent basis will give the ultimate liquidity solution much more credibility. Liquidity cost is a more appropriate starting point, even if a full-blown implementation of this approach may not be practical. Rather than look for practical shortcuts to implement risk-to-liquidity, the Committee could instead perform a similar exercise for liquidity cost.

One class of practical shortcuts should address data needs. The Committee has proposed that banks—with some assistance from supervisors—categorize broad classes of assets into a manageable set of liquidity horizon buckets. This proposal would need two simple adjustments to power the liquidity cost framework. The first would be to categorize according to bid-ask spreads rather than liquidity horizons, while maintaining the current practical approach of using relatively coarse categorizations. Bid-ask spread is one of the recommended indicators in Annex 4 to drive the liquidity horizon categorization, and so this is not a significant departure from the current proposal.

The second data adjustment would be to require some information to enable a treatment of endogenous liquidity risks. To this end, accompanying the bid-ask categorizations would also be a notion of "typical" and "large" trade sizes. Typical in this context would be the trade size for which the standard bid-ask spread would apply, while large would refer to trade sizes which would necessarily be executed at a costlier spread. From the bid-ask categories and size information, it would be possible to provide a simple rule for trade cost as a function of size and time.[3] This would require a host of assumptions, of course, but little more than the existing proposal, and with the benefit of a more realistic liquidity risk definition.

A last piece of information that is unique to the liquidity cost framework is some notion of the liquidity need for a portfolio, that is, the amount of cash the bank should be prepared to raise in a short timeframe in order to support external demands for liquidity. A broad-brush policy would be to base the liquidity cost on an assumed fixed percentage of portfolio liquidation, executed on a pro-rata basis. A more refined approach would be to make this percentage sensitive to indicators of liquidity needs, such as the use of leverage or the frequency of portfolio turnover, and to allow for an optimal choice of positions to liquidate.

Liquidity premia

The third component of the Committee's liquidity proposal concerns capital add-ons for jumps in liquidity premia. We do agree that such jumps are an important source of market risk, and should be covered by some aspect of the capital framework. We are concerned, however, that it is difficult to disentangle liquidity premia from other pricing effects. This is particularly relevant for assets that are valued using pricing models and only rarely by true market prices.

[3] We are currently conducting research in this area. We are at the Committee's disposal should they desire further details.

For these cases, which were numerous in the subprime crisis, the subjects of prudential valuation—model risk and liquidity premia—are inextricable.

We see no reason to treat liquidity premia differently from specific risk, basis risk and other sources of risk that may not be evident in the recent price history of an instrument. Our recommendation is that rather than isolating liquidity premia for a distinct capital charge, the Committee incorporates the notion of liquidity premia into an overall assessment of whether all sources of risk are covered for particular classes of instruments.

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

We agree with two of the Committee's motivations to strengthen the relationship between the standardized and internal models-based approaches: to provide a more risk sensitive regime for less sophisticated banks, and to allow for a credible fallback in case internal models are not approved at a desk level. We are concerned with the third: to establish a floor for internal models-based capital calculations.

One aspect of our concern is the operational challenge of implementing and maintaining parallel models. This can be mitigated by the definition of the standardized model: to the extent that pricing functions or data can be utilized in both models, this reduces the burden. Yet maintaining two models will still come at a cost, and banks will need to evaluate this cost against the potential benefits.

The more important aspect of our concern is that by enforcing a floor, the Committee will remove a significant incentive for banks to maintain internal models. If the cost is high and the benefits low, we fear that many banks will abandon internal models for regulatory capital, leaving most of the industry governed by the single standardized model. Even with a more risk-sensitive standardized model than those in the past, there is little to be gained from such a rigid system. We doubt that this is the Committee's intent, but we urge the Committee to be wary of this unintended consequence.

A standardized floor has already been implemented in the US, as required by the Collins Amendment to the Dodd-Frank Act. Our comments above apply to the US implementation as well.

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?

We agree with the Committee's larger objective of giving supervisors a more flexible model approval framework where they may approve or revoke models for individual desks. We believe that the proposed desk-level approach is appropriate for the Committee's objective, and do not believe that other classifications could serve instead. An advantage of a desk-level approach is how it would align the applications of models for regulatory and management purposes. A different classification, for instance by risk factors, would fail in this regard. A more serious issue is that a classification by risk factors poses serious shortcomings as an aggregation mechanism. We discuss this at length in our response to Question 6.

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?

As described in Section 4.5.2, the distinction between the two approaches hinges on how a bank is required to search for the historical period that, when used as an input into the risk model, produces the maximal measure of risk. Under the direct approach, the bank would perform the search using its full portfolio, with all relevant risk factors; under the indirect approach, to mitigate the operational burden of the direct approach, the bank would perform the search using a reduced set of risk factors and scaled down representation of the portfolio.

Common to the two approaches is the idea of choosing a single historical period on which to base the risk measures. It is this idea, and the false prudence that it provides, that is our chief concern, much more than the distinction between the two calibration approaches.

More precisely, we are concerned that by basing capital on a single period of stress, the Committee might create an incentive for a bank to hedge against only the market dynamics of that chosen period, leaving it exposed to other manifestations of market stress. Banks would be well capitalized, but only against the market dynamics of the specific period. Moreover, when the single period is applied at the desk level, it is likely that the period would be less relevant for some desks, leading to light capital requirements in those cases.

Absent any practical considerations, a response to this criticism might be that banks should test for the risk-maximizing historical period every day and for every desk, thus eliminating the incentive to hedge against only a specific set of market relationships. Conceptually, this approach is not markedly different from simply building the model on a long historical dataset, one that represents a variety of different market dynamics that could expose the vulnerabilities of a given portfolio.

As an illustration, consider a mixture model where a one-year period is selected at random from a ten-year historical dataset, a covariance structure is estimated from this period, a number of market scenarios are simulated according to this covariance, and the process is repeated

multiple times. Under such a model, it would be guaranteed that market scenarios representing the range of different stress periods in the ten-year dataset would be sampled, and each desk or portfolio's vulnerabilities exposed. The resulting Expected Shortfall would reflect those tail events that most impacted each portfolio.

The attraction of this or similar approaches is to sample an array of possible stressed market dynamics, rather than focusing on a single one. A possible criticism is that if the stress periods made up a small proportion of the overall historical dataset used for sampling, events with a large portfolio loss would be rare, and the resulting capital figure not sufficient to truly cover stress events. One response to this criticism is that the Expected Shortfall would be sensitive to the large impact events, even if they were rare. Another could be to adjust the model framework above to oversample from stress periods, while still retaining the feature of sampling from multiple historical episodes.

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?

We understand that the Committee wishes to constrain diversification benefits generally across broad categories of risk. At issue here is whether those broad categories should be based on risk factors or trading desks. We are strongly in favor of the trading desk approach for two basic reasons.

The first reason, as we alluded to in our response in Question 4, is that the trading desk level is a natural level for a bank to examine its risks for management purposes. Knowing the overall bank risk due to, for example, equity risk factors, is much less useful as a management tool than knowing the risk of, for example, the equity derivatives desk across all risk factors. We believe that internal models that are consigned to a purely regulatory use, with no other application, are doomed to be inferior to models that are used broadly across the bank.

Second, and more importantly, the application of the proposed aggregation approach to risk factors is not in fact guaranteed to achieve the constraints on diversification that the Committee desires.

For desks, the subadditivity of the ES measure guarantees that the risk of the full portfolio can be no greater than the sum of the risks of the desks. This "subadditive ceiling" establishes that the Committee's aggregation formula (Equation 3), applied with correlations set to 100%, will always bound the full portfolio risk. Consequently, applying Equation 3 with lower but still conservative correlations across desks is a reasonable way to constrain diversification benefits.

Applied to risk factors, none of the arguments of the previous paragraph apply. Putting aside risk, and considering simply the loss in a specific scenario, it is obvious that the losses on individual desks sum to the loss on the portfolio. But losses across risk factors do not always sum: the return on a foreign-denominated equity is the product (not the sum) of its equity return and foreign exchange return, and exotic products can have even richer interactions.[4] The implication of this is that not even the subadditive ceiling holds: it is not necessarily true that the sum of the risks by factors is greater than the full portfolio risk.[5] The application of Equation 3 to risk factors, therefore, regardless of the correlations chosen, does not necessarily constrain diversification, and provides only the illusion of a prudential policy.

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

Crucial to the assessment of credit risk is the time horizon. Across banking book horizons (one year or more), the nature of credit risk is that defaults can cluster across related obligors, particularly as the credit cycle heads towards its trough. Risk models adapted to these horizons must describe in some way the relationships across obligors, and between obligors and the credit cycle, that can lead to these clusters. This is the motivation behind the asset correlation structure in the Internal Ratings-Based (IRB) formula, as well as many banking book capital models. There is a natural mechanism in this structure to link default to market risks: defaults are assumed to be driven by obligor asset value processes, which can in turn be linked market factors through each obligor's equity.

At the shortest trading book horizon (ten days), defaults can occur, but more as isolated surprises, most likely accompanied by broader credit spread widening. Risk models for this horizon should prioritize the treatment of credit spread dynamics, while capturing the possibility of a small number of sudden surprise default events. In practice, most integrated models for short trading horizons do not incorporate the asset correlation structure, but rather model jump-to-default events either as independent shocks, or as linked conditionally to spread levels.

If we grant that the industry approaches to modeling at the two ends of the horizon spectrum are appropriate, the Committee is still left with the challenge of robust supervision at intermediate horizons.[6] In the absence of sufficient loss experience for direct backtesting, the

[4] This issue is developed at length in Breuer et al (2008).

[5] In fact, the literature on risk measures does not offer any upper bounds on portfolio risk as a function of risk along risk factor lines.

[6] We recall here our critique of the liquidity horizon approach, as discussed in our response Question 2. For purposes of the response to this question, we assume that such an approach, with effective risk horizons ranging from 10 days to one year, is in effect.

Committee needs at least a benchmark by which to establish reasonable capital levels under both the "default clustering" and "surprise default" situations above. This is not to suggest that these benchmarks be used as explicit capital floors; rather, the benchmarks should serve as sanity checks, with banks required to rationalize material differences between their internally modeled capital and the benchmark levels. For the default clustering dynamic, a prudential benchmark exists, in the form of the IRB model.

For the surprise default dynamic, the Committee could look to practices in risk-based margining for a simple standard approach. In this context, margin is set to cover potential losses that arise as a defaulting derivatives participant's positions are closed out. The closeout horizon is similar to the shorter trading book horizon contemplated by the Committee. For portfolios of credit derivatives, many prime brokers and central counterparties have adopted jump-to-default mechanisms, whereby margin is assessed according to a set of predefined issuer default scenarios. In a similar vein, for central counterparties to establish the size of their guarantee fund, the standard practice[7] is to evaluate the default of a number of their largest clearing members. The Committee could establish a benchmark prudential default capital level in a similar fashion, based on a fixed number of surprise defaults for a given portfolio.

As to calibration, we do not feel that it is necessary to make adjustments to the confidence level for Expected Shortfall to accommodate default risk. Indeed, one of the key advantages of Expected Shortfall is that it is sensitive to the addition of low-probability, but high-impact, events.

Finally, while we recognize the importance of recognizing default risks in the capital regime, we feel that the treatment of migration risks is unnecessary and confusing. Ultimately, the goal of the model should be to capture all risks that could impact an instrument's price. In practice, migrations tend to follow price moves rather than lead them, and as such are a misleading factor to model, especially at short horizons. We encourage the Committee to remove the discussion of migration risks as a distinct requirement of discrete credit risk modeling, and to focus instead on assuring that the model for a traded credit portfolio incorporate all potential default- and non-default-based price moves. Migration may be a convenient way to model these risks, particularly where daily credit spread data is unavailable, but in many cases a model of specific risk volatility plus the jump to default is sufficient.

We discuss the issues of model validation and backtesting in our response to Question 8.

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?

Moving to Expected Shortfall

[7] See Committee on Payment and Settlement Systems (2012).

Before discussing operational constraints, it is worth reiterating the positive qualities of ES that motivate this move. First is the sensitivity to the tail of the distribution. This is not simply a matter of ES being greater than VaR; an increase in the capital multiplier would have been a simpler way to make capital bigger. Rather, the importance of the tail sensitivity is that ES can reflect changes to the distribution that VaR may not. We have mentioned two of these already: low probability but high impact events due to defaults (in our response to Question 7), or due to one of many possible market stress periods (in our response to Question 5).

A second positive quality is subadditivity: the guarantee that the sum of the ES at the subportfolio level will be no greater than the ES of the full portfolio. It is difficult to ascertain how much practical impact this has, but as the Committee seeks to constrain diversification by aggregating desk level risks, and to add defaults or other discrete events to the loss distribution, it is important to know that there will not be any paradoxical behavior as risks are aggregated. At the same time, it is important to acknowledge that subadditivity is only relevant when applied across potential losses that are additive. As we stressed in our response to Question 6, losses by risk factor are not additive, and therefore the risks of these losses do not necessarily aggregate in any well-behaved way, regardless of the risk measure employed.

A third positive quality is stability. There are two notions of stability here. One is simulation stability—the convergence of Monte Carlo results as a bank samples from a fixed, known portfolio distribution. The second is robustness—the sensitivity of ES with respect to small changes in the distribution, such as what may occur as the input data is updated. The relative stability between VaR and ES on these points is actually a subtle issue,[8] but we expect that, on balance, the move to ES will be beneficial.

Operationally, we do not expect that many banks would have difficulty producing ES results from their existing internal risk models. In some cases, the model is parametric and the move from VaR to ES is a simple scaling. In most cases, the internal model is simulation-based (either Monte Carlo or historical), and the calculation of ES involves a slight change in how the simulation scenarios are summarized. But of course, without material changes in the models, the move from VaR to ES is not considerably different from changing the capital multiplier.

While the Committee's specification of ES as the capital statistic is unrelated to the particular form of the internal model, it is worth reflecting on the particular implications of using ES with the historical simulations method, given that this is currently the most common form of internal model used. One implication of historical simulations is that fewer scenarios are used to estimate risk (compared to what is used under Monte Carlo). This may have implications on stability.[9] Moreover, regardless of the number of scenarios, the historical simulations method produces a discrete portfolio distribution. Under discrete distributions, there is a somewhat

[8] See Acerbi (2004).
[9] In fact, this is true for VaR under historical simulations as well.

underappreciated adjustment that must be made to the traditional definition of ES in order to ensure subadditivity.[10]

Model validation

The Committee has asked for feedback on the challenges to delivering robust backtesting. To answer, we would like to address the general issue of model validation. In Section 4.2.1, the Committee specifies that the performance of the internal model will be assessed against two measures: P&L attribution and backtesting. While backtesting has been part of the internal models regime from the beginning, the explicit inclusion of P&L attribution is new. We feel that P&L attribution potentially addresses many of the shortcomings of internal models that were experienced during the crisis, and commend the Committee for adding it to the model validation framework. In fact, we believe that P&L attribution can be used even more broadly than the Committee has suggested.

As the Committee describes in Section 4.2.1, P&L attribution is intended to "determine whether the risk factors included in the desk's risk management model capture the material drivers of the bank's actual P&L". In fact, the attribution exercise should partition the actual desk P&L into three components: what is explained by the risk factors; the remainder of the constant portfolio[11] P&L; and a portion due to portfolio changes. If the third component is large, then no amount of risk factor improvement will ever explain the P&L. In this case, it is fair to conclude that the desk is strongly dependent on turnover, and therefore market liquidity, to generate profits. Referring back to our response to Question 2, this high turnover dependence would be a reason to assume a greater level of market impact exposure within our proposed liquidity framework.

The Committee's comments address mostly the first two components: the explained and unexplained pieces of the constant portfolio P&L. We agree with the Committee's view that metrics indicating how much of the P&L is explained by the risk factors should be monitored and reported regularly. We also agree that there is need for more examination of these metrics, as there are no standards in the marketplace for assessing risk models in this way. We have attached with our submission an article that examines P&L attribution on a number of fixed income index portfolios.[12]

Despite our agreement with the Committee's emphasis on P&L attribution, we should warn that the approach will not identify all potential missing sources of risk. One potential pitfall is that the method is only as good as the price sources that drive the "actual" P&L. If risk is modeled using a poor proxy (such as AAA-rated corporate yields for a AAA-rate subprime-backed CDO

[10] See Acerbi (2002) and Acerbi (2004).
[11] That is, assuming no trading, rebalancing or turnover
[12] See Finger (2012).

tranche), but that same proxy is used for valuation, then the P&L attribution will not reveal any missing sources of risk.

This leads to the more general point that some sources of risk will not be evident in the historical time series, and therefore not identified by the P&L attribution method. The potential for default is the most obvious example here: recent credit spreads on a still healthy issuer will not exhibit the jumps that would occur in the event of a credit distress, a possibility that an overall risk assessment must consider. Other potential unobserved sources of risk include the possibility that a liquidity premium or a basis jumps, even if it has been stable in recent times. Reiterating our point in our response to Question 2, we recommend that the Committee focus on all such unobserved sources of risk in a single framework instead of isolating liquidity premia for special treatment. The P&L attribution exercise should be accompanied by a qualitative assessment of other missing, but unobserved sources of risk, which in turn should be added to the model (such as in the integrated approach to credit) or covered by a capital add-on.

The second phase of model validation is backtesting. Defined narrowly, backtesting involves comparing *ex ante* forecasts of risk to *ex post* realizations of P&L. Defined more broadly, this phase of validation is to test for statistical forecast accuracy as a complement to the P&L attribution's test of model completeness. By virtue of its question, the Committee demonstrates a concern that the move to ES might compromise the efficacy of model backtesting. We feel that the greater challenge is posed by the move to the stressed calibration.

Backtesting, or indeed statistical model validation generally, is well founded when the objective of the model is clear. If the objective of the statistical forecast is to describe as well as possible the potential loss over the next day, given all information up to the present, then the traditional backtesting, as described in the Market Risk Amendment of 1996, is straightforward to extend from a test for VaR to a test on the entire forecast distribution.

On a given day, rather than simply noting whether the loss exceeds the VaR forecast, a bank could instead note the quantile of the forecast distribution corresponding to the realized loss. In other words, rather than just observing that a loss had exceeded the 99% VaR forecast, the bank would record that the loss had occurred, for instance, at the 99.3% level of the forecast distribution. Over a specified historical period, the bank could test whether its loss experience matches the distribution it forecast on each day.[13] Even more simply, on each day where VaR is exceeded, the bank could record the ratio of the realized loss to the VaR forecast. This ratio on average should be consistent with the assumed distribution.

The true problem with backtesting internal capital models is that the objective of the model is not to produce the best forecast for the next day, given all information up to the present. The goal, in contrast, is to produce a stable capital level that covers stressed market conditions (even if these are only a remote possibility at the moment) and that does not produce a procyclical

[13] See Zumbach (2007) and references therein.

capital regime. If backtesting, as described above, is the ultimate qualifier of a good model, then the best models will be those that react quickly to market conditions. Those models, of course, are the ones that produce procyclical capital—enabling in the good times, and overly restrictive as conditions turn.[14]

Paradoxically, even a stable model can produce procyclical capital if daily backtesting is used for validation. With a stable model, we would expect VaR exceedances to be rare during benign market conditions, and to cluster as markets turn volatile. A capital multiplier that is low when exceedances are low and increases if exceedances are too common would, when applied to the supposedly desirable stable risk model, produce the same undesirable dynamic capital as the reactive model.

Ultimately, we are forced to conclude that backtesting, as it is typically employed, is incompatible with the aims of a stress-calibrated capital definition. While backtesting is still a useful discipline, if only for the insight it brings from investigating a particular VaR exceedance, it should not be a primary piece of the supervisor's validation arsenal.

Our recommendation is that model validation should begin with a critical examination of P&L attribution, then the associated exercise of identifying unobserved sources of risk. Missing an important source of risk altogether, after all, is a more serious error than underestimating that risk's volatility. For statistical validation, we recommend that the Committee redouble its efforts on the specification of the stressed calibration. The Committee can help its own cause by articulating a clearer definition of this calibration, making more evident the goal of the model itself. Banks could then perform appropriate statistical validations as part of the calibration process.

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

The Committee has three objectives for the revised standardized approach: as a simple framework for banks that choose not to adopt internal models at all, as a credible fallback when a supervisor withdraws approval for an internal model at the desk level, and as a floor for the overall capital level.

For the first objective, simplicity is at a premium, and the partial risk factor approach would appear to be an appropriate direction. Our chief concern is whether this approach will in fact be simple enough, as the Committee is still quite ambitious for its ultimate risk sensitivity and granularity.

[14] See further discussion in Finger (2009).

For the second two objectives, simplicity is less a priority than consistency with internal models, and to that end the fuller risk factor approach seems more appropriate. We believe it will be difficult to reconcile all three objectives into a single standardized approach.

10. Do commenters propose any amendments to these approaches?

We share the Committee's concerns expressed at the end of Annex 7—that is, how to reflect a stress period in the calibration of either standardized approach. In fact, the notion of a standardized approach, with a single calibration applied for all banks, is at odds with the proposed calibration for internal models, where a bank should look for a period of stress tailored to its specific portfolio. One way to address this disconnect would be to move to a more unconditional approach for the internal model calibration, as described in our response to Question 5.

As to the fuller risk factor approach, the framework is appealing, and the notion of hedgeable (where exposures may be netted) and unhedgeable (where they may not) factors is a clever construct. But the true test of the approach will be in its calibration. We assume the Committee recognizes that this calibration is an ambitious undertaking, not materially different from calibrating an internal model.

As it proceeds in the calibration, we urge the Committee to bear in mind that the definition of factors, the exposure mapping to the factors, and the specification of factor volatility and correlation are all intertwined. In the example in Annex 7, we are concerned that there is an overlap of risk across the three levels of factors specified (worldwide equity index, industry equity index, price of Daimler share). Presumably, Daimler would be a constituent of the industry index, and the industry a significant component of the worldwide index, and so movements in the three would be positively correlated by construction. If the risk weight for Daimler is calibrated to the full volatility of Daimler shares, then it is effectively double counting to also include a risk weight for the worldwide index in the calculation of the overall capital for a Daimler share.

We appreciate that the structure of factors is desirable to clearly define where positions can offset one another. A better calibration that retains this structure, but provides for a more reasonable overall risk, would be to define the factors as the worldwide index return; the industry index return net of the worldwide index; and the Daimler return net of the industry index return. In this structure, it is reasonable to assert that a position in Daimler contains non-overlapping exposures to all three factors. Moreover, the Committee's proposal of treating the factors as uncorrelated is better founded under this approach.

References

Acerbi, C. (2004). Coherent Representations of Subjective Risk Aversion, in *Risk Measures for the 21st Century*, ed. G. Szego, Wiley.

Acerbi, C. and D. Tasche (2002). On the Coherence of Expected Shortfall. *Journal of Banking and Finance* **26**, 1487-1503.

Breuer, T. et al (2008). Regulatory Capital for Market and Credit Risk Interaction: Is Current Regulation Always Conservative? Deutsche Bundesbank Discussion Paper, No. 14/2008.

Committee on Payment and Settlement Systems, International Organization of Securities Commissions (2012). Principles for Financial Market Infrastructures. CPSS Publications No. 101, April.

Finger, C. (2009). VaR is from Mars, Capital is from Venus. RiskMetrics Research Monthly, April.

Finger, C. and M. Abbasi (2012). 2011 Year in Review: Risk Model Backtesting. MSCI Market Report, February.

Finger, C. and C. Acerbi (2010). The Value of Liquidity: Can It Be Measured? Working paper, June.

Zumbach, G. (2007). Backtesting Risk Methodologies from One Day to One Year. *RiskMetrics Journal*, **7**, 17-59.