MSCI

31 January 2014

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 second document on their Fundamental Review of the Trading Book. We commend the Committee on its continued development, and on its revisions to the 2012 Review that reflected feedback from the industry.

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.

Though the Review covers a number of areas, including a new standardized model and the boundary between the banking and trading books, we will focus our comments on the proposed framework for internal models. The framework in this proposal is similar to the proposal from 2012, but now with a number of aspects more precisely defined.

In choosing a specific mechanism for incorporating liquidity risk into the market risk framework, the Committee has defined a modeling paradigm with which the industry has little experience. Though we did find some aspects perplexing at first glance, we found that on further reflection, these aspects did have an interpretation that struck us as consistent with the Committee's goals.

Our comments on the new modeling framework will not recommend radical departures from the current proposal. Rather, our approach will be to illuminate some of the less obvious

implications of the new modeling framework. Our concern is that by introducing aspects that are new to the standard risk modeling paradigm, the Committee has opened the possibility that neither they nor the industry will fully appreciate how the models will behave, or what incentives they will induce, until there is significantly more experience with running the models on real portfolios. Clearly, part of the goal of the Quantitative Impact Study will be to assess the model behavior. But we feel that besides observing the models on large, complex trading book portfolios, it is also important to consider simple examples in order to build up some basic stylized facts about model behavior. Two important behaviors we identify, based on simple examples, are:

1. The liquidity penalty (that is, the additional capital required as a result of longer liquidity horizon) is smaller for assets that are strongly correlated with the portfolio.
2. The proposed method of instantaneous shocks produces an overstatement of risk for convex portfolios, relative to a method that explicitly captures a hedging strategy.

We pose the question of whether these behaviors, foreseen or not, are desirable features of the new model framework. More broadly, we encourage the Committee and the industry to consider the new modeling framework over more simple cases, in order to develop a richer understanding of other implications and incentives.

Our full comments are structured as follows. We first comment on the implications of simply using a longer holding period in the market risk model. We then interpret the mechanism to incorporate liquidity horizons, and raise a number of implications, and then do the same for the mechanism of instantaneous shocks. Finally, we provide a number of shorter comments on diversification constraints, model validation, and the stress calibration.

In closing, we applaud the Committee's efforts, and look forward to continuing the dialogue. We are available for further comment or clarification as necessary and welcome future communication.

Sincerely,

Christopher C. Finger	Carlo Acerbi	Ashley Lester
Executive Director	Executive Director	Managing Director
Applied Research	Analytics Research	Analytics Research
chris.finger@msci.com	carlo.acerbi@msci.com	ashley.lester@msci.com

1. Impact of longer risk horizons

The first implication of the new framework is that risk horizons will be longer; only large-cap equities will be treated under the current framework's 10-day holding period, while all other risk factors will be assigned a holding period of 20 days or more. Simply increasing the holding period puts greater sensitivity on a number of assumptions that are part of the risk modeling framework already.

The Committee is clear that banks should continue to model risk based on the assumption of "instantaneous shocks," that is that risk factor shocks, regardless of liquidity horizon, should flow into pricing models without accounting for any explicit passage of time. The practical attraction of this assumption is evident: it obviates the need for any assumptions about reinvestment, rebalancing or rollover.

But the instantaneous shock assumption also raises questions, even at short (10-day) risk horizons. These questions become more sensitive, however, as the risk horizon lengthens.

One question is how to handle positions with a residual maturity that is shorter than its liquidity horizon. Without an explicit treatment of this situation, it is possible that the risk model could apply 90-day risk factor shocks to a position with maturity in just a few days. This could result in a gross overstatement of risk, but it could also facilitate regulatory arbitrage: a bank could purchase an inexpensive short-term hedge, but then gain credit in the risk model for the hedge's sensitivity to larger long-term shocks to its underlying risk factor. We recommend that an explicit treatment for positions with short residual maturity be added to the framework.

A second question is how to incorporate positions with path dependency. The value of a variance swap, for instance, has no explicit dependence on the price level of its underlying, but does depend on the realized variance of the underlying as time goes by. A naïve application of the instantaneous shock framework would produce no sensitivity of the variance swap to shocks on its underlying, yet it is probably desirable that a shock intended to capture a 60- or 90-day move should influence the value of the variance swap. This implies that some degree of aging should take effect, and that some assumption on the path of such a long term shock should be made. As with the residual maturity question, this issue applies to the current framework as well, but the long horizons in the new framework make both of these more critical to address.

2. Incorporating diverse liquidity horizons

The framework introduced to incorporate market liquidity in the revised internal model-based approach proposes a simulation with varying liquidity horizons for different classes of risk factors. The horizons, of prescribed length, are meant to represent the time necessary to extinguish an exposure to a risk factor in periods of financial market stress without materially moving the price of the hedging instruments.

A first reading of the proposal arouses some questions of practical as well as fundamental nature

1. What historical dataset should be collected to generate the proposed shocks?

2. Does the method represent a simulation of real events or is it an ad hoc recipe, designed to assign different penalties to risk factors according to their market liquidity?
3. Are long-term risk methodologies too exposed to model risk – particularly regarding drift effects – to be adopted as a level playground for regulatory purposes?
4. What stylized behaviors might one expect for models constructed in this way? In particular, how will the "liquidity penalty" (the increased risk resulting from a longer liquidity horizon) depend on portfolio characteristics?

Further reflection seems to show that the proposal answers these questions, although more or less explicitly, in an unambiguous way, and that the emerging picture is robust and consistent. In what follows, we report our understanding of these questions, together with a number of important implications.

Our understanding is that the Committee is proposing the historical distribution of factor returns be built in the following way:

- A single scenario is a set of asynchronous factor returns from a common date *t* up to different dates *t+h*, where *h* is the liquidity horizon corresponding to any given factor.
- A historical distribution is obtained by recording a scenario for any past date *t* of an observation time window. Different scenarios can have overlapping returns.

This answers question 1, but triggers question 2. This dataset, because of its asynchronous nature, does not immediately appear to be a sample of meaningful events from the past.

We can reconcile the method with intuition by interpreting it as a mechanism to simulate possible evolutions of a portfolio's profit/loss over a one-year period,[1] under the assumption that any risk factor exposure is extinguished (hedged or liquidated) as soon as possible, namely at the end of the assigned liquidity horizon. This interpretation follows directly from the motivations that inspired the proposal (page 4) and from the definition of liquidity horizon (page 14).

It is when applied to a portfolio that the method generates a distribution of meaningful events, sampled from the past. The underlying assumptions may be strong, but at least the spirit of the test is clear. This provides guidance for further implementation details not yet specified by the current proposal.

Let's consider concern number 3. With the Committee's decision to move from 10-day to 1-year portfolio simulations, drift effects can no longer be neglected. But the associated model risk (drift estimation difficulty, high sensitivity of results to underlying assumptions) risks spoiling the validity of the entire proposal.

However, in the context of the above considerations, we believe that the Committee is implicitly leaving little room for model risk and in particular for artificial assumptions on drift effects. Our conclusion follows from the following considerations:

- The most straightforward implementation of the Committee proposal is a crude historical simulation based on the asynchronous empirical distribution. In such a framework, the distinction between drift and volatility components of the process is not even clearly defined.

[1] That is, the length of the longest liquidity horizon

- There is no reason why this distribution should be altered in its drift properties, for instance assuming zero mean for risk factors returns. The Committee is clearly asking to fully incorporate drift effects in the simulation, as observed in the past historical window.
- Parametric models are not ruled out entirely. For instance, one could still approximate the empirical distribution as a multivariate Gaussian, after appropriate tests. The obtained covariance matrix would have an asynchronous nature, with no model-free relationship to any synchronous one, but this would create no obstacle to a simulation engine. However, what is most important for us here, drift effects of the multivariate Gaussian would have to be estimated directly from the empirical distribution, and left unchanged.
- We conclude that the only space left for parametric models seems to be restricted to distributional approximations of the prescribed asynchronous empirical distribution. Realistically, for sake of model parsimony, the only possible choices would probably fall in the realm of elliptical distributions. No space is left for artificial assumptions on drift effects. These have to be fully taken into account by any parametric model and have to be directly estimated from the empirical distribution.

These conclusions follow from our interpretation of the proposal plus a number of implications that go beyond what is currently explicitly mentioned in the consultation paper. We invite the Committee to be more explicit with regards to the above points and most importantly to residual model risk.

We believe that altogether the proposed framework is rather robust and consistent, and is backed by a clear motivation, given by the simulation of portfolio risk exposure extinguishment over a long term horizon. Questions 1, 2 and 3 seem to have a unique consistent joint answer.

To address the fourth question, observe first that under the proposed framework, some holding periods are longer than 10 days, meaning that the volatility of many risk factors will necessarily be larger than under the current framework. Moreover, portfolios will in general comprise risk factors with a mix of liquidity horizons, meaning that the model must address, for instance, the correlation of the 20-day return on a small-cap equity with the 10-day return on a large-cap equity. This correlation can be thought of as a combination of the correlation of contemporaneous 10-day returns and the correlation of returns on non-overlapping periods. While the first of these may well be strong, the second is almost always weak, and indeed is often assumed to be zero for purposes of market risk modeling. As a result, correlations of returns in the proposed framework will be weaker than those in the current framework, unless a portfolio happens to include only risk factors with a common liquidity horizon.

There are thus two important effects to consider in the move from the current to the proposed framework: higher volatilities for risk factors with liquidity horizons greater than 10 days, which will usually increase portfolio risk, and weaker correlations within portfolios with a mixture of liquidity horizons, whose effect will depend on the makeup of the portfolio.

It is far from obvious what the combination of these effects will produce. To provide a few concrete figures, we have performed an exercise with a number of simple portfolios. In each portfolio, there is a unit exposure to the S&P 500 Index (whose liquidity horizon is 10 days), and

a unit exposure to one other risk factor (whose liquidity horizon is longer). The sign of the exposure to the second risk factor is chosen to produce either a directional portfolio (where the two positions in the portfolio are positively correlated) or a relative portfolio (where the positions are negatively correlated). We work with data from the period 2012-2013, and compute Expected Shortfall at the 97.5% confidence level using historical simulations. We first compute risk based on the current framework (with a 10-day holding period for both positions) and then based on the proposed framework (10-day holding period for the S&P 500, and a longer holding period for the second factor). In the table below, we display the "liquidity penalty", that is, the increase in risk resulting from the move from a common 10-day horizon to a full liquidity horizon approach.

				Liquidity Penalty	
First factor	Second factor	Liquidity Horizon	10-day return correlation	Directional	Relative
S&P 500 Index	Russell 2000 Index	20 days	89%	7.1%	131%
S&P 500 Index	VIX Volatility Index	20 days	-78%	9.2%	18.3%
S&P 500 Index	US Govt 10-year rate	20 days	41%	30.1%	33.9%
S&P 500 Index	CDX NA Inv Grade	60 days	-78%	33.8%	151%

One observation is that the liquidity penalty is only marginal for the first two directional portfolios; in these cases, the higher volatility from the 20-day liquidity horizon is largely offset by the weakened correlation due to the partially overlapping return periods. Where the 10-day correlation is weaker (as with the interest rate), the overall risk increase is greater. We question whether this behavior – smaller penalties for more correlated positions – is desirable from a prudential perspective.

The relative portfolios exhibit greater liquidity penalties, as the mechanism effectively removes a hedge (in this case, the negatively correlated S&P 500 position) before the end of the longer liquidity horizon. Again, we question whether this behavior is foreseen or desirable.

3. The instantaneous shock mechanism

Beyond the modeling of risk factor returns at different horizons is the implication of mixed horizons on the instantaneous shock assumption. The Committee stipulates that for a position that depends on multiple risk factors (for example, an option that depends on an underlying asset and an implied volatility), the bank should generate risk factor return scenarios appropriate to each risk factor's liquidity horizon and enter these mixed scenarios into the position's pricing model.

At first glance, this mixed instantaneous shock appears unnatural, and does not obviously describe any meaningful financial process. Upon further reflection, however, we can see that it approximates the behavior that the Committee would like to see modeled.

Recall that the Committee defines the liquidity horizon as the length of time the bank would need to close or hedge an exposure to a risk factor. Continuing with the example of the option, this definition would imply that the "true" scenario for the option position should be constructed as follows:

1. Hold the option until the end of the liquidity horizon for the underlying asset.
2. At the liquidity horizon for the underlying, delta hedge the option against moves in the underlying.
3. Hold the option and hedge position until the end of the liquidity horizon for the implied volatility.

Comparing this "true" hedge procedure to the instantaneous shock, we see that the only discrepancy comes in the period between the shorter (for the underlying) and longer (for the implied volatility) liquidity horizons. The discrepancy comes from the imperfections in the delta hedge, and can be shown to comprise one term deriving from the nonlinearities in the option with respect to the underlying, holding implied volatility constant, and a second term deriving from the change in option delta as the implied volatility changes. For simple portfolios, the first term (which we refer to as the gamma effect) dominates, meaning that we can ascertain whether the discrepancy is positive or negative based on the convexity of the portfolio. For portfolios with positive convexity, the proposed instantaneous approach generally overstates losses (and therefore risk) relative to the delta hedge procedure, while for portfolios with negative convexity, the opposite is generally true.

To assess the size of the discrepancy in methods, we examine a one-year, at-the-money straddle position on the S&P 500 Index. For this position, the underlying has a 10-day liquidity horizon, while the implied volatility has a 20-day liquidity horizon. We examine the Expected Shortfall, at 97.5% confidence level, using historical simulations on two different historical periods, the crisis period of 2008-9 and the more benign recent period of 2012-3. The results are given in the table below.

		Expected Shortfall		
Historical period	Straddle position	"True" hedging	Instantaneous	Difference
2008-2009	Long	6.5	7.9	20.5%
	Short	48.6	39.6	-18.5%
2012-2013	Long	6.5	7.1	8.5%
	Short	11.0	10.6	-3.6%

As expected, we observe that the instantaneous method produces an overstatement of risk for the long straddle position (which has positive convexity) relative to the multi-step, explicit hedging model, and an understatement for the short straddle position. Moreover, we also observe that the differences between the methods are more pronounced for the more volatile crisis period.

Our point here is not to advocate for a multi-step procedure where hedges are built explicitly into the risk model. Rather, we wish to make the observation that the proposed method, while perhaps unnatural at first, does in fact provide an approximation to a richer, more meaningful set of assumptions. Second, we point out that the discrepancy between the proposed method

and the true hedging procedure derives from a number of second order effects that are nonetheless important to understand.

4. Comments from the first review

We wish to repeat two observations from our comments on the first Fundamental Review that have not been reflected in the second version. We find it unlikely at this point that the Committee will incorporate these comments, but still feel that it is important to restate them.

First, in the treatment of liquidity, there is no charge for concentrated positions, nor indeed any dependence of the liquidity framework on position size. In the first Review, the Committee indicated that it planned to take up concentration risk through a separate charge, while in this version, there is no longer any reference to this. We continue to believe that size is an important facet of liquidity risk, and that the omission of a charge for concentration is a weak point of the new framework. One possible extension could be to provide two or three liquidity horizon tables, with shorter liquidity horizons imposed for small exposures, and longer ones for significant concentrations.

Second, for purposes of restricting diversification benefits, we feel that partitioning the portfolio by risk factors is both awkward and potentially imprudent. We reiterate that the desirable coherence property of the Expected Shortfall measure—that the sum of risks across a partition of the portfolio serves as an upper bound on the portfolio risk—is only true if the partition is a clean separation of the positions in the portfolio. A partition by risk factors, which is the basis for the proposed diversification restriction, could overlook risk factor interactions at the position level, and may not necessarily provide a strict upper bound on the portfolio risk. We find it much more natural to partition the portfolio by trading desks or business lines, which in the same framework would guarantee a prudent upper bound on portfolio risk. Moreover, we point out that by abandoning the previous approach (the restriction based on a quadratic form with floors and ceilings on risk factor correlations), the Committee has removed the need to work with the risk factor partition. The more robust diversification restriction approach offered in the new proposal could naturally accommodate a partition based on trading desks.

5. Model validation

On model validation, we reiterate our agreement with the focus on P&L attribution as a means to identify missing risk factors. The two statistics the Committee has proposed in this Review strike us as reasonable, though we would stress that there are few benchmarks for these. Before fixing strict thresholds for these statistics, we encourage the Committee not only to examine the results of the Quantitative Impact Study, but also to develop benchmark levels based on analysis of index and other representative portfolios.

The inclusion of daily backtesting is somewhat of a surprise, given the new liquidity horizon treatment and the requirement of a stress calibration. Though the Committee has proposed relatively high thresholds for backtesting exceptions, we feel it would be worthwhile to provide

more guidance on the motivation for this exercise when the model being tested (one-day holding period) is so different from the risk model being used to set capital.

6. Stress calibration

In our comments on the first Fundamental Review, we encouraged the Committee to provide clearer guidance on the definition of the stress calibration. We are pleased to see more clarity in the second Review. Moreover, we are pleased that the Committee has acknowledged the practical difficulty in obtaining long data histories for all risk factors in the portfolio. The use of the reduced portfolio in the indirect approach to obtain stressed risk figures over a long history strikes us as a reasonable compromise.

Nonetheless, we believe that basing the calibration on the worst single year for the reduced portfolio will not necessarily guarantee the desired stress on the full portfolio. The risk factors for which only a short history exists could well be those most likely to drive stress events, and yet by construction these factors will be those omitted from the long historical stress calibration. We recommend further study on how the specific mechanism for the stress calibration will behave for realistic portfolios and market conditions.