

Basel Committee on Banking Supervision Consultative Document

Fundamental Review of the Trading Book

Dated May 2012

The International Swaps and Derivatives Association, Inc.

The Global Financial Markets Association

And

The Institute of International Finance, Inc.

Further Response Covering

Measurement of Liquidity Risk

August 2013

1. Introduction

In our original response to BCBS 219 we supported the initiative of incorporating market liquidity risk into the trading book capital framework. However we also pointed out the practical problems with trying to apply a single framework across firms with different business models, position sizes and access to markets. We also outlined how firms face a trade-off between running market risk over a prolonged liquidity horizon compared with exiting positions, or hedging risk factors, in a shorter horizon but suffering the necessary costs involved of mitigating risks in that way – especially in stressed markets.

Since our initial response, the industry has invested considerable energy discussing various approaches to liquidity adjusted VaR/Expected shortfall. The purpose of this Paper is to consider the pros and cons of alternative ways of encapsulating the concern that potential illiquidity, manifesting in a stressed scenario, must be directly capitalised in a Pillar 1 regime. We remain conscious of the direction of travel by regulators in BCBS 258 to strike a balance between simplicity and complexity for the Pillar 1 measure and we believe that the liquidity horizon scaling approach proposed in BCBS 219, which we acknowledge was a proposal from industry back in 2010, may add unnecessary complexity into the framework, especially if the objective of the TBG is to provide a more coherent framework. It can also further complicate the implementation of some of the other concepts, such as granular model approval, constraining diversification benefits and enhancing the interaction between standardized and internal approaches - based on model performance on which we have provided our recommendations to the Trading Book Group previously.

We also premise our arguments on a risk factor view of liquidity. It is clear that observable market factors and hedges can often be assigned “product” granularity (for example, a credit-default swap referencing a liquid bond with clear traded volume on an authorised central counterparty). However, the positions that require hedging often blend liquid and illiquid risk factors. In attempting to characterise liquidity, it seems prudent to consider a risk-factor view, and then consider whether there is a residual prudential need to engage in a product view of liquidity.

Our argument is positioned against the perceived regulatory need to consider whether a single prescription allows one to accommodate all fair-value in a consistent way, irrespective of accounting classification. This has been characterised before as “buckets of differing liquidity”; however, we assert that this is too simplistic.

While there are several risks that characterise the fair values of different products and bank specific portfolios under different market conditions, one must assess liquidity not only in terms of outright sale of the assets but also ensure that the underlying scenario driving the realisation of losses (that stem from the changes in funding or market liquidity) is coherent.

In the extreme, if we consider local currency treasury bonds available-for-sale, funded by retail deposits, then one such scenario would be where there is a run on the bank, causing any losses on the bonds to be realised in a swift disposal of the balance sheet. The key in this particular scenario is an unexpected change in the funding position of the firm, and critically that such a scenario would typically happen only once significant losses are likely to occur to Tier 1 resources. This is incoherent with the fact that trading losses are actually typically incurred as the result of events exogenous to the overall health of the firm, whereas in other areas of fair-value, losses are incurred under scenarios that themselves are conditioned on impaired health of the firm. Therefore, when considering all fair-value assets (from trading through to non-trading assets and AFS) and the underlying funding strategy and how it characterises the potential for loss, one comes to a different conclusion on the fundamental liquidity of

the positions, and hence a single prescription of “liquidity horizon” runs the risk of describing an incoherent scenario that is not described by any probability distribution of loss.

2. Proposal

In our earlier paper we advocated evidence based intent-on-trading standard for defining the trading book with a penalty function that would move towards a less diversified and more penal capital calculation when a bank’s internal models performed in an unsatisfactory way. Implicit in this standard is the assumption that in normal times, with a high degree of confidence, risks can be hedged over a single time horizon, either by selling assets to hedge long positions, buying assets to hedge short positions, or by directly hedging the market factor sensitivities of positions. If implemented correctly, this will automatically mean that liquidity horizons in a trading book are of a short nature.

It is only in the case where a trading book is more broadly defined to capture all fair valued positions, in the trading and banking books, that multiple liquidity horizons may be required. However, this opens up a wider discussion of how market risk in the banking book should be capitalised, which we will not address here.

If regulators agree with the evidence based definition of the trading book (as proposed in BCBS219) then our proposal is to establish just one appropriate liquidity horizon that defines eligibility of products and risk factors for the book and also the capital horizon. Clearly, in order for markets to evolve, products that are less liquid but depend on at least some liquid risk factors should be allowed into the trading book but for these products the materiality of illiquid risk factors should be evidenced by the firm and appropriate capital add-ons used in Pillar 1¹. We also suggest that Pillar 2 charges be included to allow for capitalising currently liquid risk that may become illiquid in the future. This provides a convenient partition between Pillar 1 and Pillar 2 - Pillar 1 provides a capital charge based on current market conditions while Pillar 2, intended to capture risk which is not captured in Pillar 1, rightly considers the risk of loss of liquidity and how the firm would manage its positions during such market conditions. Firms will obviously have different approaches in such circumstances, depending on business model, market access, and their ability to fund through a period of stress.

3. Summary and Conclusion

There are clearly opportunities to introduce double-counts with stress-calibrations of expected shortfall models, and more broadly for fair-value risks currently in the trading book of today and likely in the non-trading book of tomorrow. The industry believes that the liquidity horizons approach may provide the opportunity to integrate the two concerns, to the extent that the overall aggregation issues can be solved.

¹ For example, assume FX options had previously only been traded to a maximum time to expiration of three years and consider the first FX option that was transacted with a time to expiration of five years. That five-year option would have one of its market factors (five year implied volatility) with a value that would have to be approximated by extrapolation from more liquid shorter tenor options. The volatility of the magnitude of five year implied vol would have to conservatively approximated as well, until the trading liquidity for FX options of five year time to expiration sufficiently developed.

The industry believes that fundamentally, portfolio hedging and diversification of idiosyncratic risks are the attributes that carve-in client activities in a capital efficient way. When that same mechanism is used to leverage up the fundamental capacity of the market and beyond, then we precisely run the risk of illiquidity in a stressed scenario. Looking at the alternatives discussed and analyzed in the appendix, it is not clear which reacts correctly as that leverage becomes unsupportable, and none of these proposals necessarily reduce pro-cyclicality.

While we continue to support the FRTB initiative and fully accept the need to adequately capture liquidity risk in the capital framework, we believe that, in the spirit of BCBS 258, the benefits of additional Pillar 1 adjustments for liquidity are more than outweighed by the costs in terms of additional complexity. Given the practical challenges of having a single prescribed measure of liquidity and the added difficulties of backtesting model performance when there are liquidity adjustments, we believe that liquidity adjustments in Pillar 1 will further complicate the already difficult problems of a more granular model approval process - constraining diversification benefits and reducing comparability between the standardized and the internal model based approach. In our previous communications with the Trading Book Group on these topics, we presented a framework with different cuts of risk (by products, risk factors, desk etc.) to fit the different purposes and most importantly, to include a weighted average of internal model approach and standardized approach based on model performance. Liquidity adjustment will take us further away from a unified framework.

We believe that liquidity risk will be adequately captured for Pillar 1 by stressed expected shortfall over a single standard liquidity horizon. Any incremental liquidity risk not captured within the stressed expected shortfall framework, which is likely to be idiosyncratic by nature, is better captured under Pillar 2. The supervisory review process should be trusted to ensure that each firm adequately identifies incremental risk and treats it in a way that is consistent with its own management strategies and its own particular market access.

Appendix

Alternatives

Is it possible to apply a satisfactory single prescription where, outside of the regulatory trading book, the amount of diversification and hedge benefit is increasingly prescribed, either through rules or through standardised calibration, and yet retain a single coherent capitalisation under Pillar 1?

If such prescription was available, then:

- The RWA framework would be indifferent to developments of accounting standards in their classification of positions and products.
- The RWA framework would be globally applicable and less subject to amendments in the localisation processes and jurisdiction specific legislation.
- One could define a Pillar 1 “Base” RWA as based on a “current” calibration of the framework, with a Pillar 1 “Counter-Cyclical Buffer” being calculated explicitly using a “stressed” calibration. Against this requirement, Tier 1, Tier 2 and Bail-In type instruments can be aligned more appropriately and transparently to the task, enhancing the quality, timeliness and applicability of disclosures.

We would further note that products that are less liquid can be capitalized under standard rules, indeed securitisations have been singled out for standard rules treatment in the revised framework. In our view the FRTB and indeed the Basel 2.5 framework have already extended Basel II in a way that to a great extent reflects liquidity risk. Expected shortfall captures ‘tail risk’ better than VaR and these tail scenarios may already reflect periods of illiquidity. Further, stressed VaR, and ES calibrated to a stressed period also reflect illiquid markets since the stressed scenarios observed at those times reflect market price movements when markets are illiquid. Finally, the Trading Book Group has previously advised that the FRTB regards the level of capital provided under Basel 2.5 as being appropriate – the objective is not to raise Pillar 1 capital further but rather to provide a more coherent framework. The QIS will provide a mechanism to ensure that approaches under FRTB deliver similar capital to Basel 2.5 and this could be done without further liquidity adjustments.

To answer the question above, we would argue that against the effects of a wide-ranging regulatory reform agenda, it may not yet be possible to apply a single prescription for measuring liquidity risk, and that Pillar 2 should still have a pivotal role. Indeed, the trading book capital regime already envisaged under the FRTB go a considerable way already to capturing liquidity risk in Pillar 1 and, given the variation across firms mentioned above, we feel that any additional capital required should be captured under Pillar 2 rather than in Pillar 1. This would avoid making Pillar 1 any more complex than it already is, while allowing firms to articulate their own liquidity management in stressed markets through a Pillar 2 charge that would still be subject to supervisory review. Pillar 2 solutions though are rarely comparable given there are few articulated standards for the regulatory assessment of firms’ own practices.

To better understand our position, consider that, without the benefit of hindsight, it is precisely due to the lack of observability of the loss of liquidity a market may endure in a stressed environment, that all approaches to this problem suffer from difficulties in calibration, and it is as important to assess each prospective solution in terms of not only the absolute output, and the ability to be conditioned on a stressed scenario, but critically also on how the metric reacts going into the next financial crisis to the

efforts of prudent risk management, the changing market vectors and the information they supply to the model.

It is equally important to reflect on the need to respect a specific firm's or market's competence with respect to the management of less liquid positions. Firms will also face different liquidity depending on the geographic zones they operate in, their access to clients and end-investors who may see stressed markets as opportunities to add to portfolios at favourable levels. Some firms will have a greater risk appetite for volatility than others and some may have capital and an ability to fund positions that means they are less likely to become forced sellers than others. If there is no regulatory appetite to respect the fact that large international firm A may be disadvantaged locally against local bank B, or vice-versa, then the expectation should be that any approach to capturing illiquidity will, necessarily, be lowest common denominator.

There are likely three rational alternatives:

- A. Divide the portfolio into risk factors by liquidity horizon, and extend the measure (VaR, Expected Shortfall (ESF)) to that horizon.
- B. Specify stressed exit costs for each risk factor, but keep a single horizon.²
- C. Calibrate a single horizon metric to a period of stressed liquidity.

Optionally, one can include a fourth approach which can be layered on top of the above alternatives:

- D. The impact of illiquidity conditioned on a stressed scenario is itself conditioned on the size of the single horizon metric: if the measure $> X$, then an additional stressed calibration is applied to determine the capital. This type of approach better addresses the concern that backtesting is weakened if in the historic prices one has never seen the effect of liquidation of the current portfolio, where the latter is precisely defined as the residual non-diversifying idiosyncratic risk³.

We look at each of these in turn.

A. Liquidity Horizons

Dividing the portfolio of risks by factor or product and aligning them to discrete "buckets" creates distortions in the overall correlation structure. It is possible to consider employing quadratic equations to correlate the overall impact of each bucket on the whole risk profile. At the product level, where liquidity resides, problems of split hedges quickly render the metric useless. Additionally, it is unlikely that any particular liquidity horizon will be appropriate for all firms and the right scaling rules will also be very difficult to ascertain. Scalar multiples of the metric can be used to affect a liquidation strategy (as used to scale 10-day VaR to a 1yr horizon), and optionally this can be adapted bucket by bucket.

² Jorion (in Value at Risk, Jorion, P. (2001), McGraw-Hill) discusses at some length an exit cost approach and we might add in passing that that analysis argues strongly that the existing 10-day VaR method should not be seen as reflecting a liquidity horizon of 10 days but rather the risk that would be taken by a firm during a longer liquidation period.

³ Such approaches can be used to create strong incentives to keep the core metrics within bounds, and the decision to break through a "level of risk" is then becomes a consideration of capital. However, in the context of VaR implementations varying in levels of complexity and the degree to which they capture the risks (RNIVs, RNIMs), this approach itself has many pros and cons.

B. Exit Costs

Exit costs permit the rules to focus on the validation and calibration of a single-horizon estimate against observable market data, and separate out the impact of potential illiquidity. This localises the calibration issue for illiquidity and introduces a floor on capital. However, careful assessment must be made to ensure the use of a floor concept creates the right incentives, again through the cycle.

Calibrating the exit costs, whether by expert judgement or based on the impact of a liquidity horizon approach on a set of representative (hypothetical) portfolios, opens up the regime to critique around applicability/suitability of those portfolios and the judgements made. However, not capturing appropriately diversification in the exit costs would be extremely penal, both in absolute terms, but given the nature of the floor in the overall risk sensitivity of the framework. Additionally, it could be very difficult to calibrate exit costs in a stressed market. This might better be considered under Pillar 2.

It is not immediately clear whether the implicit role of exit costs as a floor is the appropriate concept to then apply more broadly across fair-value in the non-trading book as defined above, and further consideration should be given to this.

C. Calibration to a Stressed Scenario

A key concern about both liquidity horizons and exit costs is precisely the loss of diversification benefit from prescribing a limited, manageable number of “liquidity buckets”. This then turns to the question of whether there is a need to explicitly reference illiquidity within the framework. It is clear that to a degree the calibration of shortfall measures to stress scenarios will capture potential illiquidity. However, there are two arguments that are poorly answered today:

- The over-reliance on historical calibration, which although internally coherent, may not actually reflect the potential for illiquidity in the future due to the structural changes occurring in the market. Advocates of addressing this issue explicitly would argue that it focuses firms on the forward-looking issue from a prudential supervision perspective.
- Assuming that potential illiquidity is implicitly caught whenever a measure is calibrated to a stressed scenario is flawed by lack of reference to whether today’s positions are indeed representative of those liquidated in that historic scenario. Again, if one has never seen the current size of position liquidated, the information is not contained within the price history, and backtesting is further weakened as a meaningful test. Without an explicit overlay mechanism such as (D) above, such an approach is open to criticism from being precisely overly dependent on the governance mechanism by which the stressed calibration is chosen.