

IRC comments

February 2009

Christopher C. Finger
chris.finger@riskmetrics.com

Amid the financial crisis of the past year, one of the many alleged villains has been the Basel II banking capital standard. Lost in the debate is the not so subtle point that the standard is still evolving, and in fact the Basel Committee is taking steps to address many of the regulation's shortcomings. One such step is the recent publication of the third consultative document on assessing incremental capital in the trading book. Though not finalized, this new regulation represents an important step to closing a gap in the existing capital rules—credit risk in trading portfolios—that has been severely exposed by recent events.

We comment here on the new proposal, from both a practical and fundamental perspective. To do so, we begin with a statement of the Committee's goals: some of which have been explicitly stated, some which we can infer from prior documents and some which we would assert should be goals, even if the Committee is not in a position to state them outright. With a set of broad principles defined, we then proceed to question whether the specific proposal achieves the goals, and to make a few modest recommendations as to how they may be better met.

A bit of history

If the one overarching theme of the first Basel Accord was consistency, that of the second accord

must be risk sensitivity, along with the acknowledgment that banks themselves are capable of measuring risks. In fact, Basel II was not the first appearance of risk sensitivity in the capital guidelines: under the 1998 Market Risk Amendment (MRA) to the original Basel Accord, banks were given the option to base part of their minimum capital requirement on their own assessment of the risk of their trading portfolio, subject to certain restrictions and standards.

For the most part, Basel II was notable for extending the internal models idea to banking portfolios. The trading book capital treatment, though updated somewhat, did not change dramatically from how it was set forth in the MRA, under which banks are assessed minimum capital in two parts: general and specific.

The general risk charge is intended to cover risks due primarily to systemic market moves; banks are permitted to use their own models to measure this risk, and capital is defined as the Value-at-Risk (at 99% confidence over a ten-day risk horizon) multiplied by a scaling factor, which is subject to supervisory discretion, but is no less than three. The specific risk charge is to cover the impact of security- or issuer-specific events: spread widening, credit migration, default, and so on. As in 1998, many risk models did not account for such events, the original specific risk charge was a standard supervisory formula. As time went on, more sophisticated banks were able to

secure approval to expand their internal model coverage to specific risk.

While Basel II left the general and specific capital charges largely unchanged, it recognized the need to better cover the default risks inherent in trading portfolios, and laid out a number of high level principles for treating incremental default risk. As the principles were presented in 2005, they were not (yet) a reaction to significant trading book losses, but rather a recognition of the amount traded credit risk had grown since the MRA.

In October 2007, the Basel Committee published a consultative document defining an Incremental Default Risk Charge. In response to both industry comments and market developments, the Committee expanded the charge to include migration risks (and consequently renamed it the Incremental Risk Charge, or IRC). The third consultative document, published in January 2009, refines the definition of the IRC, and formalizes the total trading book capital requirement as the sum of the general and specific charges from the MRA and the new IRC.

The IRC itself is defined as the VaR of the unsecured credit products in the trading book, at a 99.9% confidence level over a one-year horizon. Recognizing that imposing a one-year holding period is unrealistic for many trading positions, the IRC permits a liquidity horizon for each instrument, with the notion that at this horizon, an instrument may be rebalanced in order to maintain a constant level of risk. Other than the restriction that no instrument have a liquidity horizon shorter than three months, the definition of the model and parameters for the IRC is left largely to the banks.

A set of principles

Though tempting to jump straight to the specifics, it is important to first establish the motivations for the IRC specifically and model-based capital generally. We should begin with the very principle of risk sensitivity: required bank capital should depend on the risks of the bank. Moreover, it is probably inevitable that a risk sensitive capital regime will depend on bank internal models. Taken together, these notions are under fire, with the perceptions that through this crisis, banks were undercapitalized (a fair point) and incapable of modeling their own risks (an overreaction). These perceptions have prompted cries for a “back to basics” approach, but this is dangerous.

If by “back to basics” we mean a return to the simple capital requirements of the original Basel Accord, we should recall what motivated Basel II in the first place. The criticism of the original accord was that it was basic and inflexible, and as such produced an incentive for banks to remove low yielding (but safe) credits from their balance sheet, while retaining higher yielding (but riskier) credits. Such an incentive was no more than an academic point in a world of buy-and-hold credit, but with the rise of collateralized debt obligations (CDOs) and credit default swaps not long after, banks had a ready mechanism to act on those incentives. It is fair to say that the original development of CDOs, the descendants of which wound up at the heart of this crisis, owes to the skewed incentives of the old, basic capital regime. That is not something we should aim to go back to.

If, on the other hand, by “back to basics”, we mean a turn to more transparent capital requirements, that is a principle to hold to.

So while we should agree that a transparent, risk sensitive capital regime is still desirable, it is impossible to neglect the point that banks were undercapitalized. So the next basic principle is that we should find ways to increase capital standards, focusing in particular on the trading book, where the significant losses have occurred. At the same time, it is important to be mindful of the rules for banking book capital, as we do not want the capital standards to produce skewed incentives for where particular positions are placed.

An issue that has lurked in the background, but has come to the front in the IRC discussion is the tension between the regulatory risk horizon and the horizon at which banks manage their trading portfolios. From the regulatory point of view, the concern is the bank as a “going concern”. The aim of minimum regulatory capital, then, is that banks can withstand losses, and are able to continue their normal activities, without the need to raise additional capital. The dispute between industry and the regulators over the appropriate regulatory risk horizon, in fact, has been essentially a proxy battle over the frequency at which banks are able to raise capital. Whereas back in 2007, early industry responses to the IRC proposals argued for a shorter risk horizon, asserting that banks could raise additional capital if they suffered capital depleting losses, the events of the last year has rendered such arguments void, and it is the regulators with the upper hand in pushing for a longer horizon.

With their going concern, long horizon view, the regulators have put themselves firmly in the camp of seeking models for banking (or trading) *practices*. Regulatory capital, then, is intended to support not just possible losses on existing positions, but on all of the positions throughout the next year, as a bank

is assumed to continue its normal operations.

Internal trading risk management, on the other hand, focuses on *positions*. As such, the risk horizon is not the frequency over which the bank can raise capital, but rather the frequency at which positions can be traded. This horizon, particularly with trading portfolios and even accounting for stressed levels of liquidity, is markedly shorter than the capital horizon. It is this problem—short horizon risk on constant portfolios—that market risk managers have focused on for over a decade. And even if the models may be imperfect or incomplete, their short horizons at least allow for empirical validation.

The MRA rules represent a compromise between the regulatory and trading horizons. The regulations recognize that banks are capable of modeling short horizon risks, yet also that the purpose of capital is more than absorbing 99% worst case, ten-day losses. Thus, banks are permitted to estimate their short horizon VaR, and the regulator applies two mechanisms to convert this measure, at least qualitatively, into a long horizon one.

The first mechanism is that, while the bank calculates its VaR daily, it is assessed capital on a sixty-day moving average of the VaR measure; this adjustment is an admission that while the short term risk of a trading portfolio may indeed be quite volatile, the capital required to support longer term trading activities should be smoother. The second mechanism is the multiplier: the short horizon VaR is scaled by a regulatory factor, which is typically between three and five. The MRA rules, then, display a principle of compromise, and are a balance between quantitative risk sensitivity, as measured by the banks themselves, and qualitative prudential standards, as im-

posed by the regulator.

A consequence of this balance is that capital is not defined in precise statistical terms. Basel defines neither the horizon nor the confidence level they consider prudent, but enforce prudence through their subjective choices of multipliers. The new rules tilt this balance, attempting to define capital in a more rigorous statistical framework. And they impose on banks the task of modeling not just their positions, but their practices, over a horizon long enough to suit the regulator.

Finally, some models, as the saying goes, are useful. But with a capital regime encouraging new model developments, it is not just models that should be useful. We should also require that discussions about models be useful, and that the byproducts of these discussions, be they other models or simply heightened awareness of other issues, be useful as well.

Recipe for trouble

In the industry responses to the prior IRC consultative documents, a consistent complaint has been the lack of an overarching conceptual framework to the new rules. The confusion stems foremost from the definition of capital as the sum of the old MRA capital, based on a ten-day 99% VaR, and the new IRC capital, based on a one-year 99.9% VaR and covering a slightly different set of sources of risk. As with the MRA, there is still no precise statistical definition of capital, but the new rules are more frustrating in that they tease with the illusion of precision (all of the components are defined statistically), and yet the final aggregated number is no more than a recipe,

with no conceptual framework to justify it. What in the end does it really mean to take the sum of ten-day 99% and one-year 99.9% VaR?

While only the most ardent modelers (still) believe in precise estimates of such things as a 99.9% one-year VaR, it is not just these zealots crying for a conceptual basis for capital, the better to apply yet more complex techniques. Even the more rational among us should recognize that the lack of a conceptual framework is dangerous, and not just aesthetically bothersome. When we have no conceptual framework, capital in the end is a recipe. When a recipe is simple, such as multiplying VaR by three, it is easy to anticipate how the recipe will behave in new circumstances. When a recipe is more complicated, such as the sum of two inconsistently defined VaR figures, we do not know what we can really expect.

The examples in the paper by Breuer et al (2007) should serve as warning. Just adding the risk measures arising from different sorts of risks does not, as popular wisdom would argue, necessarily produce a conservative answer. Too many of us have gotten too accustomed to the behavior of the simplest of models, where risks add in simple ways: so we propose a recipe that under the simplest of models is guaranteed to behave as we desire, but under realistic models, we have no idea what will happen.

Industry is not blameless here either. One of the early industry counterproposals to the IRC was to add these distinct risks via the standard formula for adding standard deviations, with an “effective correlation” calibrated to match actual model results in select cases. This proposal was no better than Basel’s, in that it replaced one recipe for adding risks with no conceptual basis with another more complicated, but

equally unfounded recipe. Basel, while not coming up with something richer, was wise at least to discard this industry proposal.

What what else to do?

A better industry proposal sought to bring some consistency to the standard market risk and IRC charges. This proposal was based on the observation that three times the ten-day 99% VaR is close to the VaR for 60 days at 99.9% confidence, assuming a Gaussian distribution and the standard time scaling. Sixty days being a plausible time horizon for overall trading book risk, and 99.9% being consistent with the solvency standard initially proposed by Basel, the industry proposed using a 60-day, 99.9% VaR for all trading book capital. The important consequence of this is that it allows for a single model, encompassing all risk sources, rather than the arbitrary division of risks of the IRC, and the resulting specter of falsely conservative capital rules.

The latest proposal from Basel sticks to the separate capital charges and risk horizons, an apparent blow to the industry proposal. One rationale for the persistent one-year horizon is that Basel wishes to avoid too large a discrepancy between the capital for an asset in the banking book (which is based on a one-year horizon) and that for the same asset in the trading book (the IRC). A second rationale is clearly that Basel desires to raise capital levels by more than the sixty-day horizon would imply.

With the discussions focusing on the risk horizon and solvency standard, it is disappointing that the choice of risk statistic has never been seriously reconsid-

ered. In the consultative document, the Committee makes it clear that a shortcoming of the MRA capital rules is that they are insensitive to the sorts of rare but large events that the IRC is intended to address. Their reaction to this is to push out both the horizon and the confidence level, seemingly raising both the overall capital and the risk sensitivity of the rule.

This approach is misguided. A measure such as VaR—a specific quantile of a probability distribution—is insufficient to distinguish the risk of portfolios with differently shaped distributions. In a 1998, MRA era view, where all distributions are Gaussian (or close) the VaR is really no more than a scaled version of the standard deviation. And if all portfolios are distributed with the same shape, this is enough to differentiate more from less risky portfolios. In this setting, VaR was a sensible choice on which to base a risk sensitive capital regime.

But we are in a different setting now. If our focus is on rare but large events, then the new capital regime should be able to differentiate between two portfolios that differ only in that one is exposed to a rare but catastrophic event. As long as the likelihood of the rare event is significantly less than the VaR confidence level, the VaR of these portfolios is almost equal. One remedy is to make our confidence level stringent enough to cover what we think will be the likelihood of such rare events. But this is a fragile regime, and only risk sensitive in the particular range of probabilities that are defined a priori. A different statistic—in particular expected shortfall, or the expected loss, conditional on experiencing a loss in the tail—produces a measure that is at once sensitive to the risks we desire to measure, and not as sensitive to the arbitrary choice of confidence level.

Constant level of frustration most definitely useful.

By far the most controversial element of the Basel proposals has been the notion of an assumed constant level of risk. The constant level of risk and its companion notion, the liquidity horizon, are in a sense concessions by the Committee that a one-year risk horizon with fixed positions is unrealistic for most trading portfolios. This author has been torn on the proposal, in fact praising it in an earlier column.¹

To rationalize somewhat, our praise at the time was for the constant level of risk notion as the first explicit acknowledgement that with long risk horizons, it is imperative to address the problem of portfolio aging. The assumption of constant positions, while adequate for short trading horizons, is unrealistic in the context of the “going concern” outlook that Basel aims to apply for regulatory capital. We maintain that Basel is due praise for promoting what is the first serious industry discussion of this topic.

Praise is due as well for the emphasis on the liquidity horizon. While its role in this particular modeling framework is debatable, the assessment of a liquidity horizon for trading positions is a worthy goal in its own right.² That the liquidity horizon for trading positions is a quantity that can materially impact capital requirements is a guarantee that these will be the subject of significant scrutiny and debate. Regardless of the ultimate application of the liquidity horizon, this debate will necessarily raise awareness and promote transparency of liquidity risks. While not a stated goal of the IRC proposal, this transparency of liquidity is a desirable side benefit, and discussions thereof

The constant level of risk framework, on the other hand, is an invitation to arbitrary recipes, and for the more cynical, to manipulation of the minimum capital through unneeded complexity. In a utopian world, we would build a risk model that simulated each asset through time. For each asset, at the appropriate liquidity horizon, we would sell the asset, realizing a gain or loss, and reinvest the proceeds. The reinvestment would be into a new hypothetical asset with the same characteristics as the original position at initiation. The motivation for such a rebalancing strategy is that it is expected to reduce risk relative to the strategy of holding all positions for an entire year.

But the effects of the rebalancing strategy are subtle, and it is unclear when the strategy actually does reduce risk. Among other things, the outcome depends on three features of the utopian simulation:

- Structure of credit (or risk) curves. It is a common that an issuers’ credit curve is upward sloping, that is, that if a default event is to occur, it is more likely to occur in the long term than in the near future. If this is true for most issuers, then rebalancing reduces our risk by exposing us to a succession of “near futures” rather than a single long term.
- Inter-period correlations. In efficient markets, we can typically assume (and validate) that market moves in one period are uncorrelated with market moves in the next. But with ratings migrations and defaults, it is much more difficult to argue market efficiency, and ef-

¹See Finger (2008).

²Indeed, in a parallel consultative process, Basel has issued a set of principles for sound liquidity risk management.

fects such as ratings momentum are well documented.

- **Risk dynamics.** It is acknowledged that both default probabilities and volatility vary with time. Thus, it is not certain that the risk of a BBB position in six months, after my initial BBB position has downgraded, will in fact be a constant level of risk.

In liquid markets, where risk factors evolve continuously (or approximately so), it is plausible that we might defend such a set of assumptions. But ratings migrations and defaults do not evolve continuously, and what is more, are subject to very human, discrete processes: when does a firm choose to file for bankruptcy, or an analyst to review a rating. When migrations and default are the risk factors, clarifying a model with rebalancing will always come down to some articles of faith. And when capital depends in complex ways on articles of faith, it is naive not to believe that faith could well be corrupted.

It is naive as well to believe practical a utopian simulation with asset-by-asset rebalancing; any implementation of an IRC model will be a mere approximation to this ideal. At least in the special case where there is a single liquidity horizon, the consensus emerging in the industry is to obtain the one-year distribution by summing successive short-horizon portfolio losses: for a liquidity horizon of three months, we define the three month portfolio loss distribution, then sample four times from this. Embedded in this operation is the assumption that in our succession of “near futures”, each period is independent of the next. Mathematically, this procedure amounts to computing a convolution of the short horizon loss distribution.

The convolution approach, while tractable and seemingly consistent with the spirit of the IRC definitions, is problematic. Most importantly, the approach does not define a model in terms of a process; we have not specified how the portfolio evolves from one period to the next. As such, there is no natural way to generalize the approach to a portfolio with a mix of liquidity horizons, a situation sure to arise in practice. Consider a portfolio composed of three-month and six-month liquidity horizon positions. Applying the convolution to either portion of the portfolio is straightforward, but how do we make sense of the dependence between the six month loss experience on the less liquid subportfolio and the month four to month six loss experience (including hypothetical rebalanced positions) of the more liquid portfolio?

In the end, the convolution approach, which appears to be the only tractable way to address the constant level of risk proposal, is in truth no more than a scaling rule, albeit an opaque and complicated one.

A humble proposal

Acknowledging the basic principles and goals of the IRC proposal, we offer a recommendation for a slight modification. We have in a sense a cap on what the IRC should be, in the form of the Basel II internal ratings-based (IRB) capital charge—the capital that would be required were the assets in question to sit in the banking book. In the limit where our trading positions are highly illiquid, these positions should be treated exactly as banking book positions.

We should establish a limit on the other end of the spectrum: what the IRC should be for very liquid

positions. Basel has established that they view three months as a minimal liquidity horizon, and so we should define this liquid-IRC based on a three-month holding period. While it is appealing to establish a risk measure (VaR) and confidence level (99.9%) to be consistent with the IRB charge, we are faced with the shortcomings of the VaR measure discussed previously. It would be sensible, then, to base the definition on an expected shortfall, with confidence level set to match the prudential standard of the IRB charge, at least in some special cases. The Committee should reserve the right to apply a capital multiplier in this case, in order to raise overall capital levels without compromising the desirable risk sensitivity of the measure.

With our two limits established, it remains to specify how capital should be computed for portfolios of an arbitrary liquidity horizon. Here, we must admit that the constant level of risk approach, while appealing in spirit, is unattractive in practice, and will in all likelihood lead to obfuscation without any real modeling advances. The liquidity horizon, on the other hand, is a worthwhile endeavor on its own. We propose, then, that banks establish and defend liquidity horizons for all positions, subject to a prudential regulatory floor, and calculate the average liquidity horizon across the portfolio. Moreover, the liquidity horizon definition should play a role in the needed definition of “illiquid positions” within the specific IRC rules.

The final IRC charge should then be obtained by a simple interpolation rule (not necessarily linear, but still transparent) between the two limiting cases,

based on the portfolio average liquidity horizon. Granted this proposal is itself but a recipe, and not a coherent statistical definition of solvency. But in the spirit of the old MRA compromise, it is a recipe that makes fully transparent the dependence of capital on the liquidity horizon, while at the same time preserving an appropriate level of capital for the most and least liquid portfolios. Most importantly, it is a recipe that encourages the industry to focus its research efforts on items that are valuable in their own right—short-term default probabilities, liquidity horizons, basis risks—and limits the most onerous operational requirements and least valuable modeling discussions at the same time.

Further reading

- Basel Committee on Banking Supervision (2008). Principles for sound liquidity risk management and supervision. September.
- Basel Committee on Banking Supervision (2009). Guidelines for computing capital for incremental risk in the trading book. Consultative document. January.
- Breuer, T. et al (2007). Regulatory capital for market and credit risk interaction: Is current regulation always conservative? Conference on the Interaction of Market and Credit Risk, Berlin, December.
- Finger, C. (2008). The once holy grail, RiskMetrics Research Monthly, January.