RiskMetrics Group

Research Monthly

The once holy grail

January 2008

Christopher C. Finger
chris.finger@riskmetrics.com

I recently participated in a conference on the interaction of credit and market risk. For the most part, it was very enjoyable, with interesting talks, stimulating discussions and a chance to visit Berlin. But I also felt a sense of disappointment, not with any of the talks specifically, but with the idea that there was a need for such a conference at all.

Ten years ago, we were fresh from the release of the CreditMetrics technical document, and our thinking turned to the greater modeling challenge that lay ahead: the integration of credit and market risk. There was a somewhat simplistic call to arms in a magazine piece titled "The Do's and Don't's of CreditMetrics" where the author urged "*Don't* add the risks from RiskMetrics and CreditMetrics." It was a daunting problem, but we figured that this would be what we would focus on for the next several years. That there would be a conference on the topic ten years later, with the problem still unsolved, would have been a discouraging thought back then.

So what happened? Or to put it more bluntly, why have we, the risk community, seemingly failed at presumably our great challenge of the last decade?

Why apart?

To put some perspective on the problem, before we consider how to put market and credit risk together, we should ask ourselves how they came to be apart in the first place. There are two reasons: one good, and one sensible, but ultimately bad.

The good reason is that the modeling challenges with market and credit risk are different. The study of market risk has to do with the fluctuation in value of an asset or portfolio, both forecasting the magnitude and identifying the sources of the potential fluctuation. For most assets, we are well served by historical observations: the past is a good indication of magnitude of future market moves, and of future relationships between risk factors and risky assets. Surely, plenty of disclaimers apply here, but historical data is what we have, and more often than not, it is of some use. Partly, this is true because market fluctuations are more or less continuous and thus chances are, we will have observed *something* in the past, and that something will at least look somewhat like what could happen in the future.

With credit risk—the risk that a counterparty or issuer defaults on its obligations—the past, at least in the manner we use it for market risk, is of no use at all. One year of daily returns on a foreign exchange rate may tell us something about how that rate may move next week, but knowing that a counterparty has met all of its obligations in the past does not lead us to an estimate of its probability of defaulting in the future. Credit, fundamentally, is a risk of something happening that has not happened before, but that we

acknowledge as a possibility. Models, then, need to somehow construct defaults from proxy or extrapolation, rather from direct empirical evidence.

The bad reason for the separation of risks is cultural. Due in part to the accounting standards of the 1980s, and cemented by the 1998 Basel Market Risk Amendment, banks discretely assign positions to either the banking or trading book. Traditionally, the banking book, containing mostly loans to be held until maturity, was the realm of credit risk: market fluctuations were unimportant since the loans would never be traded, leaving defaults as the sole source of risk. The trading book, on the other hand, was the realm of market risk, with trading fluctuations the source of risk, and default events either irrelevant or extremely unlikely. So it was natural that market risk models came to be associated with the trading book and credit risk models with the banking book.

While the separation of assets into banking and trading may be a proper partition, the distinction of market and credit risks is not. Even prior to the boom in credit trading, there was credit risk in the trading book, in that derivative counterparties could default, and market risk in the banking book, if only in the form of foreign currency denominated loans.

If we consider the total economic capital required to support a bank's aggregate risk-taking activities (both its banking and trading books), the real problem is how to integrate the risk of two portfolios with somewhat overlapping sources of risk. Unfortunately, the problem came to be viewed as that of integrating risks, which is the easy part, and because portfolios stopped being part of the problem statement, the more difficult, but more important issues were never addressed, or for that matter even raised.

The real problem – horizon

So what are the issues? To greatly simplify, there are three: modeling credit risk in trading book portfolios, modeling market risk in banking book portfolios and modeling the total risk of the combined portfolio. But this statement of the problem leaves us in the (somewhat arbitrary) language of bank capital. A much better statement of the problem is to use a word we I have not mentioned yet: horizon.

More than anything, what distinguishes the two books are their respective risk horizons. The trading book allows for decisions to be made intra-day, or daily, or weekly—frequencies such that risk measures at a five- or ten-day horizon are the most meaningful. The banking book, consisting of assets to be held to maturity, is characterized by a longer, if somewhat ill-defined horizon. And importantly, by framing the problem in terms of risk horizon, we transcend the banking conventions, and operate in a framework that is applicable to any financial institution, not just those that fall under the Basel auspices.

When considering the risk horizon, we notice one other fundamental difference. At least to first order, the key risk parameter for market risk, volatility, scales with the square root of the risk horizon; the key parameter for credit risk, default probability, by contrast scales linearly. The implication is that the relative importance of these two risks changes with the risk horizon we consider. Armed with a perfect joint model of the two risks, we might not care about this statement, but in a world where practicality is in order, it is a profound point, and suggests that the modeling choices we make should be governed by the horizon we consider. The solution depends on

the problem we are trying to solve. It is not one holy grail we are searching for, but many.

Traded credit risk

With the growth of traded credit over the last decade, it could be argued that the relative importance of the short horizon problem has grown. What the industry has done well is to incorporate spread risk into the standard market risk frameworks, with spreads acting simply as another market risk factor.

But credit spreads are not just another market risk factor; indeed, the reason they exist in the first place is the possibility of a default event, though we do not see such an event in the spread's historical data. It is common to assert that a model for spreads that admits the possibility of large widening events is sufficient to describe default risk in traded portfolios. But defaults are more than a really bad spread widening: they are a bad spread widening punctuated by a definitive loss, with no possibility of a move back to a benign level. In other words, there is no mean reversion when you experience a default.

This modeling flaw rears its head not just in risk modeling, but in the marketing of credit structured products, in particular Constant Proportion Default Obligations (CPDOs).[1] At any point in time, such a structure is composed of a deposit account plus a leveraged position on a credit index. The leverage is determined by a formula, and is generally high when the credit index spread is high. Thus, the structure will take on more credit risk as spreads widen, and profit as spreads tighten again (or mean revert).

[1] See O'Kane (2007).

Early versions of CPDOs promised Libor plus 200bp and still received AAA ratings. Such a risk-return profile should make us scratch our heads. If the modeling of the spread dynamics to price the structure treats spreads as a mean-reverting factor, then certainly a structure that depends on spread mean reversion will look appealing. But spreads can only mean revert if defaults do not occur. Failing to account for default as an absorbing state means the pricing is essentially conditional on no defaults occurring, and it should be no surprise that any credit structure is deceptively attractive.

So our first big question is whether today's state of the art—modeling spread risk as a standard market risk factor—truly captures the credit risk in short-horizon portfolios. One argument is that with horizons this short, the likelihood of a default occurring is quite small; so even if we modeled defaults well, it would have little impact on our standard risk measures. Just for perspective, consider that the historical annual default rate for BBB-rated firms is about 20bp, implying the two-week default rate is less than a single basis point. Even for a portfolio containing one hundred such firms, there is still less than a one percent chance that a single default occurs within the risk horizon, and little material influence on, say, the 95% Value-at-Risk.

So how do we assess the risk of these rare but severe default events? Without doing any modeling, but just acknowledging that defaults can occur, we can perform a jump-to-default stress test: we suppose that all the obligations of a single firm (or a group of firms) suddenly jump to their recovery value, and evaluate the impact on our portfolio

through all exposure—from simple bond positions to structured portfolio credit derivatives—to the firm (or firms) under stress. Though perhaps trivial on the surface, jump-to-default stress tests serve as an important complement to standard risk measures, in that they can expose sensitivities to large events in portfolios that are by design hedged against incremental spread changes.

To go further than jump-to-default stresses, we do need to propose a model for default events in a short horizon. But this is pointless if we do not also begin to monitor statistical trading risk measures—expected shortfall, for one—that are sensitive to tail events, as default events in any two week period are.

The other flavor of default risk in a trading portfolio (any trading portfolio, not just one explicitly trading credit) is counterparty risk. This is another area where the market's evolution has radically changed the modeling problem. Ten years ago, the challenges in estimating potential future counterparty exposure were in projecting market rates over long risk horizons (horizons comparable to the life of the derivatives in question) and, particularly in the aftermath of the Asian crisis of 1997, in assessing the relationship between long term market moves and the credit health of counterparties. Crucially, counterparty exposure was a long horizon problem.

In the past decade, developments in credit mitigation techniques, particularly those that provide for dynamic updates to collateral or margin, have changed the modeling questions. With agreements in place to update collateral on a frequent basis, depending on mark-to-market changes, the risk begins with the question of objective (and frequent) mark-to-market and mark-to-model valuations. For future exposure, the horizons are only as long as the worst case in the update cycle, often no longer than a week. And the true difficulties are in describing the operations of the credit mitigation, and in assessing the risk that these procedures may not function as advertised, a risk that falls into the operational domain.

Beyond ten days

So is the horizon ever long? Certainly, we still need to worry about market risk in banking books; but beyond banking, and the arbitrary distinction between risk types, are a large class of investors who simply have a long risk horizon. Both pensions and insurers manage assets with risk horizons of one year or more. Moreover, neither pensions nor insurers are burdened by the cultural division of market and credit risk; what is important is total return over a relatively long horizon, and the impact of this return on the institution's solvency. The horizon is what characterizes the problem, not the risk type.

One important implication of the long horizon is that both credit and market risk matter. So we have to do the hard work both of forecasting our market factors at a horizon longer than we are accustomed, and of integrating the default events with the market factors. But here, there has been progress on both fronts. Zumbach (2007) and references therein discuss our thoughts on the market factors, and the structural (or Merton) framework gives us a natural link between market factors and defaults. Granted, there are details to fill in, but a framework exists, and some of

[2] Kupiec (2007), for instance

the work presented in Berlin[2] illustrates how the various pieces can be put together.

On the technical modeling front, then, we have plenty to show for the last ten years. There is, however, a more mundane implication of a long horizon that has generated disappointingly little discussion. With a one-year horizon, we have to assume something about what happens to our portfolio over the course of the year. This involves simple (essentially static) things, like what we do with matured positions or coupon payments, and complex (or dynamic) things, like what assumptions we make about closing out a position that starts to lose money. For lack of a better term, I will use the term "portfolio aging" to refer generically to all of these aspects. Our assumptions about portfolio aging all matter, as they combine to essentially define the risk we are measuring. We may have an adequate hammer, but we are further away from finding the right nails.

Basel returns

Default risk in the trading arena has not escaped the scope of Basel 2 regulation, and a new consultative paper was released in October to elicit comments on how to include incremental default risk into trading book capital calculations. The banking industry has already responded in part to this paper: Smillie and Epperlein (2007) summarize the issues nicely, and make a number of counterproposals of their own.

Discussion of this issue is only beginning (implementation of this specific capital calculation will not be required until 2010), but the consultative paper sets an important precedent by explicitly acknowledging the portfolio aging problem. The paper proposes that default risk be calculated at a risk horizon of one year, but admits that a bank would not hold a liquid position for a full year if the position were deteriorating. Thus, the "buy and hold" approach to the portfolio used in traditional credit modeling would overstate risk. To compensate for this, the paper introduces the notion of *Constant level of risk*: the assumption that each position is rebalanced at a particular liquidity horizon back to its original riskiness.

But while the discussion of portfolio aging is welcome, other aspects of the proposals in the consultative paper, as well as those of Smillie and Epperlein, are less attractive. One aspect is the continued reliance on Value-at-Risk for the capital statistic. Critiques of VaR are often actually critiques of a particular risk estimation technique or modeling assumption, or of the general notion of forecasting risk from historical information. Here, my focus is narrow: the use of a percentile of the forecasted distribution as the relevant risk statistic. As mentioned before, when dealing with default risk for short horizon portfolios, we are faced with the prospect that defaults have no material impact on typical risk measures, such as VaR at a 95% confidence level. The consultative paper, seemingly to address this difficulty, proposes that credit risk on the trading portfolio be measured with a one-year horizon and a 99.9% confidence level. Market risk would still be measured at ten days and 99% confidence. The portfolio dynamics discussion is to mitigate the anticipated criticisms that the one year holding period is not realistic.

Smillie and Epperlein propose a compromise: that both credit and market risk be measured at a 60-day horizon at 99.9% confidence. The desire for consis-

tency is commendable, but the solution is worse than the problem. By moving market risk measurement to a 99.9% confidence, we would forfeit the most attractive feature of the current market risk rules: that the VaR used for capital can be empirically validated.

Thus, the move toward long horizons and high confidence levels is well intentioned but artificial. The good intentions are to find a set of parameters where VaR is actually sensitive to default risk. Unfortunately, this leads us to horizons that are unrealistic for trading portfolios, and confidence levels at which we cannot hope to do empirical validation. These artificial parameters, in the end, result from our choice of VaR as a risk measure. The bottom line is that the real business problem—dealing with rare but large losses—calls for something other than a percentile of a risk distribution.

Come together

More problematic is the continued separation of market and credit risk. All of the discussion—from the risk horizon for credit risk to how credit and market losses should be correlated—implicitly assumes that we are still operating in a segregated world where market-related losses live in one neighborhood and credit losses in another.

One of the Berlin papers, Breuer et al (2007), argues that simply adding credit and market risk capital numbers is not, as conventional wisdom would have it, "conservative". The conventional wisdom rests on the assumption that credit- and market-related profit and loss are additive. The authors discuss foreign-denominated loans, where the obligor's repayment ability is denominated in their home currency. A loss on the loan results from a combination of weakening of the home currency (a market risk phenomenon), coupled with the obligor's repayment ability (a credit risk factor) not keeping up with his then higher foreign obligation. But where an actual loss occurs, the natural definition of market-related loss (assuming the currency moves but the obligor's repayment ability does not) might not produce a loss at all; similarly, the natural definition of the credit-related loss might also be zero. So the sum of the market- and credit-related losses may be zero while the real loss is positive. No "conservative" correlation assumptions suffice to compensate for this interaction.

The contortions that the market/credit segregation require do not end in the banking book. For a corporate bond valued at 100 today, we may have a market risk model that forecasts prices down to 80 due to interest rate and spread movements and a credit risk model that works with the probability that the bond defaults and goes to a recovery value of 60. If the issuers is seen to deteriorate, the future default loss gets anticipated in the price, and this anticipated loss is reflected in our spread risk model. So what is the potential credit loss? 40? 20? The only correct response is not to answer the question, and to work with a single profit-and-loss for which there is no need to define market- and credit-related portions.

Derivative counterparty risks raise similar issues. Suppose we enter an interest rate swap today with a present value of zero. For a standard credit risk model, we estimate a potential future credit exposure and worry about the counterparty defaulting, in which case we lose this exposure amount. But we can only lose from a default if we have first had some

appreciation of the contract due to market moves; we cannot simultaneously lose money due to both market- and credit-related events. This is not a statistical relationship for which we should be arguing over whether to apply a -40% or-80% correlation: above all, we should be modeling our overall profit-and-loss jointly, not in two distinct frameworks. Eventually, I do need to attribute risk, or more crucially assign bonuses, to the trader and the credit analyst, but not before making sure the overall profit-and-loss and risk are correct.

The next ten years (or less)

So we are still faced with challenges in the interactions of credit and market risk, but these are different challenges from what we anticipated ten years ago.

The first challenge is in how we define risk, and has to do with breaking the artificial distinctions between credit and market risks. Faced with the intractable problems of foreign-denominated loans, corporate bonds or derivatives with counterparty risk, the consultative paper skirts the issue by using the word "incremental" in its definition of default risk. But that just leaves us with the problem of what incremental default risk really is. More productive than arguing this point would be to state what total risk is, and worry that the total capital we calculate is adequate. If we get that right, the distinction of credit and market risks becomes much less important.

The second challenge is how we express the impact of credit risk at short horizons, where we recognize that the risk exists but also that even if we model the risk correctly, it will not materially impact our standard risk measures. This argues for transparency in jump-to-default stress tests, and for the use of risk measures, such as expected shortfall, that are sensitive in rare but material losses.

The third challenge is portfolio aging. Or perhaps the third challenge is to give a better name to this challenge, and the fourth to solve it. To even define risk at long horizons, we need to be clear about what we assume happens during the horizon. It is time for more open discussion of this issue. Watch this space.

Further reading

- Basel Committee on Banking Supervision (2007). Guidelines for computing capital for incremental default risk in the trading book. Consultative document. October.
- 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.
- Kupiec, P. (2007). An integrated structural model for portfolio market and credit risk. Conference on the Interaction of Market and Credit Risk, Berlin, December.
- O'Kane, D. (2007). Recent developments in the credit derivatives market. Presentation at EDHEC Alternative Investments Conference, London, November.
- Smillie, A. and Epperlein, E. (2007). Compliance by default. *Risk*, **20**(12): 96–100.
- Zumbach, G. (2007). Backtesting risk methodologies from one day to one year. *RiskMetrics Journal*, **7**(1): 17–60.