A Flexible Analytical Method to Calculate the Specific Risk Surcharge

Jan Kwiatkowski, Riccardo Rebonato and Lorenzo Liesch
RBS

January 13, 2007

Abstract

We present a simple analytic method to estimate the specific risk surcharge and an 'implied' method to calibrate it. Our method borrows from the conceptual foundations laid out in the Basel II approach, accounts for concentration effects, satisfies the comaprability requirement with the banking book treatment laid out in the Basel/IOSCO document and is easy to implement. The central idea is that, instead of calculating the conditional expected loss, we estimate the conditional percentile loss. We do so without having to make use of a Monte Carlo simulation.

We suggest a calibration methodology and present emprical data.

1 Background

Since the late 1980s/early 1990s regulators have granted banks the ability to calculate the market risk component of their regulatory capital for trading-book positions using internal models. The charge for general market risk was extended to include 'specific risk', ie, loosely speaking, the risk arising from the possibility that an issuer may be downgraded or may default.

Throughout the second half of the 1990s the regulators then devoted much effort to bringing up-to-date the regulatory capital rules pertaining to the banking book. This complex enterprise is normally referred to as the Basel II project.

As the Basel II enterprise was coming to a conclusion, the regulators came to realize that the trading-book treatment of defaultable bond and related positions was less satisfactory. Originally, banks tended to have on the trading book high-quality bond portfolios, for which default was a very rare occurrence. As credit derivatives, CDOs and CDO tranches became more widespread on the trading book, however, it was no longer reasonable to assume that default could be safely ignored. At the same time it was recognized that VaR, the cornerstone of the internal-model approach, is poorly suited to dealing with default events.

In order to 'clean up' the loose strands left over by the Basel I and Basel II project, the Basel/IOSCO Committee focussed under the 'unfinished business' rubric with the so-called 'specific risk surcharge'. This strand specifically addressed the topic of how the possibility of issuer default should be translated into a capital charge for *trading book positions*.

The Basel/IOSCO Committee's document of July 2005 requires that a bank must "capture in its regulatory capital, default risk of its trading book positions..." and that the bank must demonstrate that its approach "meets a soundness standard comparable to that of the internal-ratings-based approach for credit risk as set forth in the Revised Framework, under the assumption of a constant level of risk ....". This is reflected in the September 2005 amendment to EU Directive [2000/12/EC].

The exact interpretation of the words "a soundness standard comparable" has not been fully clarified, and a dialogue between the industry and the regulators is currently under way to give precise and operational meaning to this expression. The exact interpretation is very important: if comparability were interpreted as identity, some impact studies carried out by internationally active banks indicate that the market-risk capital charge could increase fivefold.

We propose in this paper a methodology to address the issue of the trading-book specific risk surcharge that has several desirable features:

1. it rests on the conceptual foundations laid down by the Basel II approach;
2. it does provide "a soundness standard comparable to that of the internal-ratings-based approach";
3. it can be extended from one to many factors; if the multi-factor approach is taken, banks that already use any of the existing multi-factor models can easily adapt their model inputs to be able to apply the method we propose;
4. it is simple to implement and to understand, computationally non-prohibitive, and uses the technology that most bank will already have put in place for their Basel II calculations;
5. it makes use of probabilities and percentiles that are not so remote or high, respectively, as to be of dubious meaning or interpretation;
6. thanks to accurate approximation techniques that have recently been introduced, it does not require a simulation, and can be implemented analytically.

We would like to stress at the outset the importance of point 3. above: since many banks have already opted for a multi-factor approach form their banking book calculations, we show how to transform their model-dependent parameters for use with our methodology. This feature should considerably enhance its appeal, comparability and ease of use.

We explain the approach by starting from a very brief review of the conceptual framework that our approach shares with the Basel II (banking book) approach.

2 Overview of the Approach

In order to calculate the specific risk surcharge in a manner compatible with the Basel Committee requirements one can obviously take very different approaches. A conceptually appealing way to systematize this array of alternatives is to consider two possible 'extreme' positions and a continuum of intermediate solutions. At one extreme we have a direct 'transliteration' of the Internal-Ratings-Based (IRB): in this approach one assumes that exactly the same underlying factor affects all the obligors. One then stresses the 1-year probability of default to the desired percentile (say, 99.9^{th}), and one calculates the resulting expected loss, conditional on the common factor affecting the defaults being at this high percentile level. At the opposite extreme one could assume complete independence of the obligors, construct the associated (unconditional) loss distribution, and read off the resulting curve the desired percentile. We argue that a more desirable solution can be found somewhere in between these extremes, by using a hybrid approach: we create an appropriate conditional loss distribution and from it estimate a (conditional) percentile that gives rise, thanks to the conditioning, to the desired *unconditional* percentile. This can be seen more precisely as follows.

The capital calculation under the IRB approach for the Banking Book is equivalent to finding the expected (in the statistical sense) one-year loss under stressed probabilities of default. In the banking book the unstressed expected loss is subtracted from this. However, in the trading book losses are incurred relative to the mark-to-market values, which already implicitly allow for (unstressed) expected losses, so this adjustment does not apply. Also, in the banking-book treatment there is a maturity adjustment to allow for the fact that assets may be on the banking book for more than one year. By the short-term nature of the holding of trading-book instruments this adjustment does not apply either. With these caveats, if one wanted to apply the same approach to trading-book defaultable assets, the stressing of the probability of default for any asset/obligor could be viewed in two stages:

1. A 'correlation', R, is first computed using the IRB formula based on the asset/obligor's risk rating (and hence probability of default):

$$R = 0.12\left[\frac{1-\exp\left(-50PD\right)}{1-\exp\left(-50\right)}\right] + 0.24\left[\frac{\exp\left(-50PD\right)}{1-\exp\left(-50\right)}\right] \tag{1}$$

   Equation (1) is a logarithmic interpolation between a correlation of 12% and a correlation of 24% with respect to the probability of default. See Basel Committee on Banking Supervision, 2003, para 241, and Ong (2005). R has an interpretation as the square of the correlation between the asset/obligor and a single systematic 'factor'. This systematic factor can be thought of as representing 'economic conditions' insofar as they affect corporate defaults.

2. The correlation R is then fed into a formula[1] (the Vasicek (1991) formula)

[1] ibid. This formula is described below.

that uses the single factor model to compute the stressed probability of default, which is equivalent to the probability of default under (in this case) 99.9^{th} percentile (1-in-1000-year) severe conditions.

Since in this approach the capital is based on stressed *expected* losses, the total capital requirement would be obtained simply by summing the requirements for individual assets/obligors. Indeed, it is because in the IRB approach one looks at the expected loss that, despite the obligors being not independent, one can still sum their individual contributions in a way that 'knows about' the common correlation. Indeed, the expectation is the only common statistics enjoying this simple additivity property. Of course, the price to be paid for this simplification is the inability to capture concentration effects.

Our reasons for not applying the IRB-based approach directly to issuer default risk in the trading book are threefold. First of all, the high turnover of trading-book assets means that a 1-in-1000-year scenario, representing systematic effects, is unrealistic and unsuitable for these assets. Second, we do not think that such high percentiles can be meaningfully estimated or validated. Given the inextricable link between percentile level, sampling frequency and time homogeneity of the phenomenon being studied, we are not even sure that speaking of a one-in-a-thousand year event in a financial context is meaningful. Third, the method takes no account of concentrations of risk, either at issuer level or at industrial sector level.[2]

At the other extreme from the IRB approach, we could ignore any systematic effects and treat all defaults as being independent. We could use, say, a Monte-Carlo simulation to compute the required unconditional percentile (in this case the 99.9th) of 1-year default losses. We will refer to this as the 'idiosyncratic' approach. This would automatically account for issuer-concentration effects: pursuing the example in the footnote above of two exposures to AA names, the loss distribution arising from two positions in a AA name of \$100m each is now very different from the loss distribution from one position of \$1m in a AA bond and one position of \$199m in another AA bond. However, we consider that even for 'liquid' trading-book positions completely ignoring co-movements among obligors is not satisfactory, and some account should be taken of systematic effects.

To overcome both shortcomings (insensitivity to concentration in one case, and to systematic effects in the other) we propose a method which is a hybrid between the IRB and idiosyncratic approaches. The rationale is as follows.

In the context of the trading book, we wish to take into account *systematic* effects that are less extreme than the 1-in-1000-year effects that are used in the IRB approach for the banking book, while retaining the same level of overall 'conservatism'. Suppose we stress the PDs according to a 1-in-n-year systematic

[2] This directly follows from the assumptions underpinning the IRB approach (which assumes an infinitely-granular portfolio): once the N-obligor problem is decomposed into the sum of N stand-alone contributions, the expected loss arising from, say, two positions in a AA name of \$100m each is the same as the expected loss from one position of \$1m in a AA bond and one position of \$199m in another AA bond.

effect (below we suggest $n = 8$) . This corresponds to the $100(1-\frac{1}{n})^{th}$ percentile of the systematic loss factor. Conditional on this event, the defaults are independent and we can easily calculate the conditional loss distribution. We now want to choose some percentile, say q, of this conditional distribution that is consistent with the IRB methodology. In particular we would like our method to converge approximately to the IRB methodology in the extreme case, $n = 1000$. The latter is equivalent to the mean of the 1-in-1000-year distribution. To the extent that the mean is close to the median the IRB charge will approximate the 99.95^{th} percentile of the overall (stressed) distribution. When $n < 1000$, we wish to achieve an overall (combined systematic and idiosyncratic) charge equivalent to the 99.95^{th} percentile. Given the independence of the systematic and idiosyncratic effects, we require

$$\frac{1}{n}\left(1-\frac{q}{100}\right) = 1 - 99.95\% \tag{2}$$

ie

$$q = 100\,(1 - 0.05\% * n) \tag{3}$$

This gives the systematic percentile $p = 99.9$ and the idiosyncratic percentile $q = 50$ (the median) for the extreme case $n = 1000$, as required. With $n = 8$, as proposed, we get $p = 87.5$ and $q = 99.6$.

As regards systematic effects, we want to take into account potential credit-cycle effects. Specifically, we want to take into account the systematic effect induced by the worst year in an average credit cycle. Taking 8 years as the average length of a credit cycle, we take $n = 8$ to capture a 1-in-8 year systematic effect (i.e. the 87.5^{th} percentile, or $p = 87.5$). Note that we 'simply' have to estimate the default probabilities in the downturn of a 'typical' cycle. We do not have to place ourselves at the bottom of such an exceptionally adverse cycle as to obtain the 99.9^{th} overall (ie, over cycles and over years in the cycle) percentile. Since economic cycles have periods of five to ten years, our approach makes the statistical estimation problem at least feasible and realistic. Conditional on our being in this *worst* year in a *typical* cycle, we use the IRB methodology to compute the 87.5^{th} percentile stressed probabilities of default according to the Vasicek formula, using the IRB formula referred-to above for computing the correlations.

Conditional on being in the worst year in the typical cycle, and given our (and the IRB) one-factor approach, all the obligors are now independent[3]. It is at this point that we depart from the IRB approach. Rather than taking the conditional mean loss (as in the IRB approach), we apply the idiosyncratic approach using the stressed probabilities of default, and estimate an appropriate percentile of the simulated losses over a given time horizon. Here we have taken a 1-year horizon. However, we consider that a shorter (say, 3-month) time horizon would be more than adequate in view of the greater liquidity of the trading book. This is where the interpretation of the comparability requirement mentioned in the introduction comes into play: we argue that what ultimately

[3] We explain later how this strong assumption can be relaxed.

matters is to obtain the same remoteness from default (the same percentile) given the time frame appropriate to each environment. This time frame may well be one year in the banking-book case, but is certainly much shorter for trading book exposures.

Whatever the 'correct' choice of the horizon may be deemed to be, by our approach we achieve the following:

1. we have taken into account systematic effects due to the credit cycle;
2. we have taken into account concentration effects;
3. we can adjust to whatever time horizon and overall percentile are considered appropriate.

We are justified in using the idiosyncratic approach with the stressed probabilities of default, because under the one-factor model the individual defaults are conditionally independent (ie, the stressing of probabilities of default accounts for all dependency between issuers)[4].

In practice most banks' credit models will be multi-factor models that include country/industrial sector effects, but again it should be a simple enough matter to adapt such models in such a way that dependencies are conditioned on a given percentile of a single global factor[5]. In the final section of this paper we illustrate a very basic industrial-sector model that we have constructed from market data.

As for the probabilities of default, as a first step one can use standard rating-based probabilities of default. We regard this as being extremely conservative in the trading book context.[6]

3 Technical Background

In this section we give a brief summary of the basic theory on which most one-factor models (including the IRB approach) are founded.

Consider an issuer/obligor k with rating r_k, and corresponding long-term, unconditional probability of default, $PD^T(r_k)$, over some horizon, T. We define a corresponding default boundary as

$$d_k^T \equiv \Phi^{-1}\left[PD^T(r_k)\right] \tag{4}$$

[4] There are, of course, the usual complications of dependencies between parent company and subsidiary company defaults, etc; however, by adjusting the input probabilities of default and setting factor-related correlations to zero, banks should be able easily enough to adapt their existing credit risk models for these purposes.

[5] The term 'global' is not uniquely defined, but could be, for example, the main principal component of the full factor structure

[6] Work performed by Moody's KMV (see Reference 2.) strongly suggests that the occurrence of 'surprise' defaults, (i.e. those not anticipated by their ratings) are relatively rare. The implication is that for a trading book, where the portfolio can be, and typically is, rebalanced to maintain a stable level of risk, the standard probabilities of default greatly exaggerate the risk. One could easily reduce the probabilities of default accordingly.

where $\Phi^{-1}[\cdot]$ is the standard normal inverse cumulative distribution function. We stress that, given the definition of $PD^T(r_k)$, d_k^T is the long-term (unconditional) average of the default boundary. To lighten notation, in the following we shall often simply write d_k and $PD(r_k)$, but the dependence on the horizon should be kept in mind. Obviously, if one ascribes to each obligor in a given rating the same probability of default, as one does with one-factor approaches, then all obligors with the same rating have the same default boundary.

We also define a standard normal random variable (SNRV), $\widetilde{y}_k$, representing the 'standardized assets' of an issuer/obligor of rating r_k. If we take the issuer as defaulting if the standardized assets have fallen below the default boundary, d_k, by time T, ie, when $\widetilde{y}_k < d_k$, we can write:

$$\Pr[\widetilde{y}_k < d_k] = \Phi[d_k] = \Phi\left[\Phi^{-1}[PD(r_k)]\right] = PD(r_k) \tag{5}$$

as required.

Under the one-factor model that underlies the IRB approach, the standardized assets for different issuers/obligors are correlated through their dependence on a common SNRV, $\widetilde{\epsilon}$, say, thus:

$$\widetilde{y}_k = \rho(r_k)\widetilde{\epsilon} + \sqrt{1-\rho^2(r_k)}\widetilde{x}_k \tag{6}$$

where $\widetilde{x}_k$ is an obligor-specific 'idiosyncratic' SNRV, independent of $\widetilde{\epsilon}$ and of all other obligor-specific variables. We note in passing for future reference that, by construction and for any $\rho(r_k)$, $\widetilde{y}_k$ is therefore also a SNRV. The quantity $\rho(r_k)$, it is a 'correlation' (equivalent to the square root of R as described in Section 2 above) which, under the IRB approach, is computed as a given function of $PD(r_k)$. As for $\widetilde{\epsilon}$, we can think of this quantity as a 'systematic' economic effect that influences all issuers/obligors according to Equation 6.

For any given systematic effect, say, $\underline{e}$, we can compute $\Pr[\widetilde{y}_k < d_k|\widetilde{\epsilon} = \underline{e}]$, ie, the *conditional* probability of default given $\widetilde{\epsilon} = \underline{e}$, as follows. From Equation 6

$$(\widetilde{y}_k < d_k) \Longleftrightarrow \left[\left(\rho(r_k)\widetilde{\epsilon} + \sqrt{1-\rho^2(r_k)}\widetilde{x}_k\right) < d_k\right] \Longleftrightarrow \tag{7}$$

$$\Longleftrightarrow \widetilde{x}_k < \frac{d_k - \rho(r_k)\widetilde{\epsilon}}{\sqrt{1-\rho^2(r_k)}} \tag{8}$$

Since $\widetilde{x}_k$ is a SNRV independent of $\widetilde{\epsilon}$ we can also write:

$$\Pr[\widetilde{y}_k < d_k|\widetilde{\epsilon} = \underline{e}] = \Phi[d_k(\underline{e})] \tag{9}$$

where

$$d_k(\underline{e}) = \frac{d_k}{\sqrt{1-\rho^2(r_k)}} - \frac{\rho(r_k)\underline{e}}{\sqrt{1-\rho^2(r_k)}} \tag{10}$$

The quantity $d_k(\underline{e})$ can therefore be thought of as the *conditional* default boundary, which, as intuitively it should, moves 'closer' under 'bad' economic conditions (ie, as e decreases).

4 Banking Book Treatment

The IRB capital calculation for the banking book utilizes the probability of default conditional on a 1-in-1000-year (adverse) economic effect; this is achieved by setting in Equation (10)

$$\underline{e} = \Phi^{-1}[1 - 99.9\%] = -\Phi^{-1}[99.9\%] \simeq -3.09 \tag{11}$$

By direct substitution this gives the 'Vasićek' formula for the conditional probability of default:

$$PD(r_k|\underline{e} = -3.09) = \Phi\left[\frac{\Phi^{-1}(PD(r_k))}{\sqrt{1-\rho^2(r_k)}} + 3.09\frac{\rho(r_k)}{\sqrt{1-\rho^2(r_k)}}\right] \tag{12}$$

Based on this conditional probability of default, the regulatory capital for the obligor is then calculated as the corresponding adjusted[7] *conditional expected loss*. This, in fact, is just the product of the conditional probability of default and the 'default delta' (i.e. the expected monetary loss given default) of the obligor. This, in the banking book, is subject to a maturity adjustment for exposures of maturity greater than 1 year.

Note that, since the calculation is based on expected loss, the capital requirements for different exposures can be computed independently, and there is no allowance for concentration effects. As we argued above, we consider such an adjustment inappropriate for the trading book.

5 Trading-Book Treatment (Basic)

As explained in the introductory section, we have adapted the methodology to allow for the short-time-horizon nature of the trading book and for concentration effects. The underlying principles are that we condition on more realistic systematic effects, $\underline{e}$, and that instead of computing the *conditional expected loss*, we compute the appropriate *conditional percentile*. More specifically, we take $\underline{e} = \Phi^{-1}\left[\frac{1}{n}\right]$, representing a 1-in-$n$-year adverse scenario. (We choose n equal to 8, corresponding to a 8-year cycle, in which case $\underline{e} = \Phi^{-1}[1 - 87.5\%] = -1.15$).

To compute the desired percentile, we proceed in four stages: first, we use the IRB formula (Equation 1) to determine the correlations[8]. Second, with these correlations and using Equations (9) and (10) we compute a stressed probability of default for each issuer:

$$PD(r_k|\underline{e} = -1.15) = \Phi\left[\frac{\Phi^{-1}(PD(r_k))}{\sqrt{1-\rho^2(r_k)}} + 1.15\frac{\rho(r_k)}{\sqrt{1-\rho^2(r_k)}}\right] \tag{13}$$

[7] for maturity and *unconditional* expected losses, as explained in the introductory section.

[8] The assumption is made here that the correlation is a function of the unconditional rating, ie, of the unconditional probability of default.

Third, from these conditional probabilities of default we must create the conditional distribution of losses. One could do so in a single Monte Carlo simulation. It is more efficient, however, to use the analytic approximation of Andersen, Sidenius and Basu (2003) to estimate the conditional distribution of losses (Jaeckel (2005), offers some suggestions as to how the procedure can be improved).

Lastly, from the distribution we can directly estimate the desired (ie, in our case the 99.6^{th}) percentile.

6 Trading Book Treatment (Multi-Factor)

In this section we suggest how our hybrid approach can be extended to multi-factor models. We also show how banks that use such models can adapt the model inputs to be able to apply our method. In the following section we will give an example of a simple application of this method that we have used for our corporate portfolios. The underlying model is similar to the KMV and Credit Metrics approaches, but we have used our own analysis for estimating the required parameters.

For each issuer, in addition to its rating, r_k, we denote its industrial sector code by s_k. An issuer's standardized assets, $\widetilde{y}_k$, are therefore modelled as being driven by a systematic sector m-vector, $\widetilde{\eta}$, and its own idiosyncratic effect $\widetilde{x}_k$, thus:

$$\widetilde{y}_k = \rho(r_k)\left(\alpha_k^T \widetilde{\eta}\right) + \sqrt{1-\rho^2(r_k)}\widetilde{x}_k \tag{14}$$

where

- the components $\widetilde{\eta}_j$, $j = 1, 2, ..., m$, of the vector $\widetilde{\eta}$ are SNRVs, independent of the idiosyncratic factors $\widetilde{x}_k$;
- the components $\widetilde{\eta}_j$ are correlated via a real-symmetric matrix $C(\eta)$;
- α_k is a normalized vector of weights, with components α_{kj} that describes the 'loadings' of the kth issuer on the industrial sectors. The weights α_{kj} are such that
$$\alpha_k^T C(\eta)\alpha_k = 1 \tag{15}$$

The components $\widetilde{\eta}_j$ ($j = 1, 2, \ldots m$) of the vector $\widetilde{\eta}$ represents *systematic* effects, corresponding, for instance, to industrial sectors, country of business, etc. Each of these components is assumed to be a function of the single economy-wide factor, $\widetilde{\epsilon}$, and of an idiosyncratic component:

$$\eta_j = \theta_j \widetilde{\epsilon} + \widetilde{\zeta}_j \tag{16}$$

where the quantities θ_j are constant coefficients, and the $\widetilde{\zeta}_j$ are independent of $\widetilde{\epsilon}$ and distributed according to $\mathcal{N}(0, \sqrt{1-\theta_j^2})$. For future reference, we note i) that this immediately implies that each component $\widetilde{\eta}_j$ is an SRNV and ii) that, since

$\widetilde{\epsilon}$ and $\widetilde{\eta}_j$ have unit-variance (by definition and by construction, respectively), θ_j is simply the correlation between $\widetilde{\epsilon}$ and $\widetilde{\eta}_j$.

As for the sector-specific components, ζ_j, they need not be independent. We write their covariance matrix as $C(\zeta)$. We note in passing, and we show in the Appendix, that for any realization $\underline{e}$ of the economy-wide factor $\widetilde{\epsilon}$, and, in particular, for any required percentile of $\widetilde{\epsilon}$, the model conditional on $\widetilde{\epsilon} = \underline{e}$ can be formulated exactly in the same form as Equation (14) with adjusted parameters.

If we take the correlation structure among the systematic factors, $\widetilde{\eta}_j$, as 'primitive', (ie, given by econometric estimation), we need to specify the quantities θ_j and $C(\zeta)$ in such a way as to preserve the exogenous correlation structure among the $\widetilde{\eta}_j$s. Clearly there is no unique solution to this. Ultimately, it comes down to choosing a model for $\widetilde{\epsilon}$ to proxy the 'global economy' as a function of the chosen factors $\widetilde{\eta}_j$. We pursue an implied approach, by fitting a simple model using historical Credit Default Swap (CDS) spreads. How this can be done is shown in the next section.

7 Implied Calibration: A Multi-Factor Example

The underlying ideas are as follows. The general multi-factor model in Equation (14) allows for the generic issuer to be affected by m economy-wide factors, $\eta_j, j = 1, 2, ..., m$, which represent systematic effects corresponding to, say, industrial or regional sectors. In order to carry out an implied calibration, we restrict the generality of the approach by mapping each issuer into one and only one such sector. Different issuers may, of course, be mapped onto different sectors. It is in this sense that the approach is still multi-factor.

The proposed calibration approach is not as restrictive as it may seem if we interpret *the* one factor onto which each issuer is loaded as some linear combination of industrial, regional and other factors. If the linear combination were chosen to mimic, say, the first principal or independent component, the single-loading picture could still provide a good explanation of the drivers for the chosen issuer. With the implied approach we do not have to make the identification explicit, but the interpretation can give some reassurance as to the scope and potential of the method.

Let the sector of issuer k be s_k. Equation (14) now becomes:

$$\widetilde{y}_k = \rho(r_k)\widetilde{\eta}_{s_k} + \sqrt{1 - \rho^2(r_k)}\widetilde{x}_k \tag{17}$$

where in the expression $\alpha_k^T \eta$ a single weight is picked out by our mapping assumption. Clearly, the requirement $\alpha_k^T C(\eta)\alpha_k = 1$ implies that

$$\alpha_{kj} = \delta_{j,s_k} \tag{18}$$

and Equation (17) immediately follows. As for Equation (16), it can be written as:

$$\eta_s = \theta_s \widetilde{\epsilon} + \widetilde{\zeta}_s \quad \text{for any sector } s \tag{19}$$

Our goal is to determine the θ coefficients (the correlations between the sectors and the economy) and $C(\zeta)$, the covariance matrix of the 'residuals', ζs, in a way consistent with the exogenous correlation matrix $C(\eta)$. In the implied approach, we have no a priori model for the θs or for $C(\zeta)$; rather we fit these quantity directly from time series data. In order to estimate the sector-economy correlation coefficients (the quantities θ_s), and the cross-sector correlations of the 'residuals' (the quantities ζ_{ss}) following the implied route, we have used a database of historical values, by S&P rating and maturity, of CDS spreads for the market as a whole and broken down by sectors. We consider this much more satisfactory than using equity prices, which may contain a lot of 'noise' not related to credit. Notwithstanding that the PDs inferred from spreads of CDs (the so-called 'risk-neutral' PDs) contain a risk premium that generally make them higher than historical PDs, we believe that hey give us the best chance of estimating the required correlations. The fact that the general levels and volatilities of the implied PDs may be different from the historical ones does not in itself affect the correlation estimates. By extension we would expect the sample correlations of the corresponding implied quantities $\eta_s(t)$ to reflect the correlation structure of the quantities η_s.

Consider rating r with default probability $PD(r)$ and default boundary $d(r) = \Phi^{-1}[PD(r)]$. For each sector, s, we observe every day the historical ten-year constant-maturity credit spreads[9]. From these we back out a time series of constant-maturity *implied* probabilities of default, $PD_s^T(r,t)$: these are default probabilities as seen from time t out to horizon T, for (any) issuer of rating r in sector s:

$$PD_s^T(r,1), PD_s^T(r,2), ..., PD_s^T(r,t) \tag{20}$$

Again to lighten notation, when no ambiguity can arise the superscript T will be dropped in the following. It is important to highlight the link between the CDS-implied quantities and the unconditional probability of default introduced in Equations (4) and (5). The former reflect not only the market's perception of the probabilities of default as seen at time t, but also the market's aversion to default risk. The latter are estimated directly from default histories as representing long-run "through-the-cycle" probabilities of default. In a risk-neutral world we would expect the long-run averages over time of the probabilities $PD_s^T(r,t)$ to be close to the imputed probabilities $PD^T(r)$ for all ratings r and sectors s. In practice the presence of market risk aversion means that the former tend to be higher than the latter.

According to the model, these PDs vary because of changes in the underlying economic factor, $\widetilde{\epsilon}$, and the sector-specific systematic effects, $\widetilde{\varsigma}_s$. At each point in time they are in fact the probabilities of default under the prevailing economic conditions. From the time series of the CDS-implied default probabilities we now compute the time series of prevailing default boundaries for sector s:

$$d_s(r,t) = \Phi^{-1}[PD_s(r,t)] \quad \text{for } t = 1, 2, ... \tag{21}$$

[9]We chose the liquid spreads for CDSs whose maturity approximately match the average length of credit cycles (8 years in our case).

Now from Equation (17), (dropping the k subscript), for any issuer with rating r in sector s:

$$\widetilde{y} < d(r) \Longleftrightarrow \rho(r)\widetilde{\eta}_s(t) + \sqrt{1-\rho^2(r)}\widetilde{x} < d(r) \tag{22}$$

so given that $\widetilde{\eta}_s = \eta_s(t)$, the value of the systematic sector effect at any time t, we have[10]:

$$\widetilde{y} < d(r) \Longleftrightarrow \widetilde{x} < \frac{d(r) - \rho(r)\eta_s(t)}{\sqrt{1-\rho^2(r)}} \tag{23}$$

So the conditional probability of default at time t is:

$$PD_s(r,t) = \Phi\left[\frac{d(r) - \rho(r)\eta_s(t)}{\sqrt{1-\rho^2(r)}}\right] \tag{24}$$

By equating this with the observed *implied* probability of default, $PD_s(r,t)$, for each time $t = 1, 2, ...n$ we can back out a time series of implied values of the sector factors:

$$\eta_s(t) = \frac{d(r) - \Phi^{-1}\left[PD_s(r,t)\right]\sqrt{1-\rho^2(r)}}{\rho(r)} \tag{25}$$

In order to be able to estimate the θs from Equation (19) we now require a proxy for the time series of the global economic factor, $\epsilon(t)$, for $t = 1, 2, ...n$. The way we have chosen to do this is to use the spread data for 'global' (i.e. all-sector) indices (distinguished by rating), to back out implied 'global' probabilities of default, $\overline{PD}(r,t)$ for $t = 1, 2, ...n$ and for all ratings r. We assume that the indices of spreads are sufficiently diversified (across sectors) as to reflect global economic effects only; this is equivalent to assuming that the implied probabilities of default reflect global effects only, so that the appropriate model is the single-factor model of Equation (6). Using equations (9) and (10), we see that, conditional on $\widetilde{\epsilon} = \epsilon(t)$, the probability of default of an issuer with rating r is:

$$PD(r|\widetilde{\epsilon} = \epsilon(t)) = \Phi\left[\frac{d(r) - \rho(r)\epsilon_s(t)}{\sqrt{1-\rho^2(r)}}\right] \tag{26}$$

By equating this with the observed implied probability of default, $\overline{PD}(r,t)$, for each time $t = 1, 2, ...n$ we can back out a time series of implied values of the global economic factor:

$$\epsilon(t) = \frac{d(r) - \Phi^{-1}\left[\overline{PD}(r,t)\right]\sqrt{1-\rho^2(r)}}{\rho(r)} \tag{27}$$

Therefore, thanks to the observations below Equation (16), the quantity θ_s can be interpreted as the correlation between $\widetilde{\eta}_s$ and $\widetilde{\epsilon}$. See also Equation (19). We

[10] In the following equation η is *not* a random variable - hence no superscript $\widetilde{}$. The same applies to the next equation.

therefore estimate the parameter θ_s as the sample correlation between the time series $\eta_s(t)$ and $\epsilon(t)$.

Finally, we have to estimate the cross-sector correlations among the sector-specific effects, $\widetilde{\varsigma}_s$. These can be obtained by direct calculation of the sample correlation. Alternatively, we can back out the correlation among the residuals, $\widetilde{\zeta}$, analytically. Recall in fact that the standardized sector indices $\widetilde{\eta}_k$ and $\widetilde{\eta}_j$ can be defined in terms of market index plus residual $\widetilde{\zeta}$ as:

$$\widetilde{\eta}_k = \theta_k \widetilde{\epsilon} + \widetilde{\zeta}_k \tag{28}$$

$$\widetilde{\eta}_j = \theta_j \widetilde{\epsilon} + \widetilde{\zeta}_j \tag{29}$$

Therefore, using the orthogonality between the residuals and the market, the correlation $Corr(\zeta_k, \zeta_j)$ can be derived from the other known terms as:

$$Corr(\widetilde{\zeta}_k, \widetilde{\zeta}_j) = \frac{Corr(\widetilde{\eta}_k, \widetilde{\eta}_j) - \theta_j \theta_k}{\left(\sqrt{1-\theta_k^2}\right)\left(\sqrt{1-\theta_j^2}\right)} \tag{30}$$

This formally concludes the calibration of our model.

8 Empirical Performance of the Model

If the implied probabilities of default behaved according to the model

- 1. we should get the same implied systematic sector effects, $\eta_s(t)$, from the data for different ratings and different maturities; and
  2. for a long enough time series the quantities $\eta_s(t)$ should be approximately normally distributed with mean 0 and standard deviation 1 for all ratings, sectors and maturities.

For our analysis we have made use of approximately five years' worth of data (1,141 dates), for seven maturities (6m, 1y, 2y, 3y, 5y, 7y and 10y) and seven ratings (AAA, AA, A, BBB, BB, B, C) for several sectors. We present the analysis for three typical sectors, ie, financials, energy and industrials, for all ratings and for all maturities. We empirically observe the following. Looking at prediction 1, we observe that

- for any given maturity, the averages and standard deviations of the quantities $\eta_s(t)$ display a strong similarity across sectors (see Tabs I and II);
- for all sectors, all the standard deviations of the quantities $\eta_s(t)$ are very constant across maturities; the averages, however, display more variation (see Figs 1 to 6);
- the averages of the $\eta_s(t)$ are not the same across ratings, but the standard deviations do display strong similarity for ratings from AAA to BB (see Figs 7 and 8).

The empirical results are clearly not consistent with total constancy across ratings and maturities. However, the observed variability is orders of magnitude smaller than, say, the variability of the underlying credit spreads, indicating that the transformation from the spreads to the quantities $\eta_s(t)$ does capture some of the required features.

Moving to the second prediction (ie, that the quantities $\eta_s(t)$ should be SNRVs), the same tables and figures clearly show that the estimated $\eta_s(t)$ tend to have standard deviations on average lower than 1, and averages lower than 0. Also, Fig 9, which refers to the energy sector (similar consideration apply to the other sectors), shows that high-rated issuers tend to have 'too much' volatility, and lower-rated ones 'too little'. Given the simplicity of the model and the total absence of any phenomenological 'free parameters' in the procedure, it is however reassuring and encouraging that an agreement is achieved between theory and predictions within a factor of two.

Finally, we stress that our estimates of the quantities θ_s, which are *correlations*, are unaffected by means and standard deviations.

Tab I: Averages of the $\eta_s(t)$ across sectors for different maturities averaged over ratings

Tab II: Standard deviation of the $\eta_s(t)$ across sectors for different maturities averaged over ratings

Fig 1: Averages of the quantity $\eta_s(t)$ across ratings for different maturities (Financials)

Fig 2: Averages of the quantity $\eta_s(t)$ across ratings for different maturities (Industrials)

Fig 3: Averages of the quantity $\eta_s(t)$ across ratings for different maturities (Energy)

Fig 4: Standard deviation of the quantity $\eta_s(t)$ across ratings for different maturities (Financials)

Fig 5: Standard deviation of the quantity $\eta_s(t)$ across ratings for different maturities (Industrials)

Fig 6: Standard deviation of the quantity $\eta_s(t)$ across ratings for different maturities (Energy)

Fig 7: Averages of the quantity $\eta_s(t)$ across maturities

Fig 8: Standard deviations of the quantity $\eta_s(t)$ across maturities

Fig 9: Time series for the quantity $\eta_s(t)$ for the energy sector

9 Conclusions

We have presented a simple and flexible method to calculate the specific risk surcharge for defaultable trading-book positions which is conceptually similar to the approach proposed in the Basel II, Pillar I framework. We modify the standard procedure in such a way as to allow for concentration effects. The idea is to reach the desired loss percentile by 'splitting' the required move in the normalized assets between a systematic and an idiosyncratic component. This can be done both within a single-factor and a multi-factor framework. We have presented an implied calibration methodology for the multi-factor method.

Empirical tests are clearly not perfectly in agreement with the theory, but the degree of congruence between predictions and results is encouraging.

10 Appendix

The model represented by Equation (16) implies that, conditional on $\widetilde{\epsilon} = \underline{e}$, the quantities η_j in Equation (14) have mean $\theta_j \underline{e}$, variance $\left(1 - \theta_j^2\right)$, and covariance matrix $C(\zeta)$. Using Equation (16) we can re-write equation (14) as:

$$\widetilde{y}_k - \underline{e}\rho(r_k)\left(\alpha_k^T \theta\right) = \rho(r_k) \sum_{j=1}^{m} \alpha_{kj} \widetilde{\zeta}_j + \sqrt{1 - \rho^2(r_k)}\widetilde{x}_k \tag{31}$$

The standard deviation, S_k, of the right-hand side of this expression is:

$$S_k = \sqrt{\rho^2(r_k)\left(\alpha_k^T C(\zeta)\alpha_k\right) + \left[1 - \rho^2(r_k)\right]} \tag{32}$$

It is clear therefore that we can write Equation (31) as:

$$\widetilde{y}_k^* = \rho^*(r_k) \sum_{j=1}^{m} \alpha_{kj}^* \widetilde{\zeta}_j + \sqrt{1 - \left[\rho^*(r_k)\right]^2}\widetilde{x}_k \tag{33}$$

with:

$$\left[\rho^*(r_k)\right]^2 = 1 - \frac{1 - \rho^2(r_k)}{S_k^2} \tag{34}$$

$$\alpha_{kj}^* = \alpha_{kj} \frac{\rho(r_k)}{\rho^*(r_k) S_k} \tag{35}$$

and

$$\widetilde{y}_k^* = \frac{\widetilde{y}_k - \underline{e}\rho(r_k)\left(\alpha_k^T \theta\right)}{S_k} \tag{36}$$

This has the same form as Equation (14), and is (conditionally) a SNRV. The conditional problem can therefore be analyzed by the same methodology as the unconditional problem with the appropriate transformation of parameters. Specifically, the stressed PD, $\Pr\left[\widetilde{y}_k < d_k | \widetilde{\epsilon} = \underline{e}\right]$ can be written as

$$\Pr\left[\widetilde{y}_k < \frac{d_k}{S_k} - \underline{e}\frac{\rho\left(r_k\right)\left(\alpha_k^T \theta\right)}{S_k}\right] = \Phi\left(\frac{d_k}{S_k} - \underline{e}\frac{\rho\left(r_k\right)\left(\alpha_k^T \theta\right)}{S_k}\right) \tag{37}$$

Acknowledgement 1 *It is a pleasure to acknowledge useful discussions with Dr Mark Joshi, who suggested the split between systematic and idiosyncratic risk used in this paper.*

References

Andersen L, Sidenius J, Basu, S, (2003) "All your hedges in one basket", Risk, November, 67-72.

Basel Committee on Banking Supervision,(2003) Consultative Document, The New Basel Capital Accord, April 2003

Bluhm, C, Overbeck L, (2005), "Explaining the Correlation in Basel II: Derivation and Evaluation", Chapter 2 in "The Basel Handbook", Michael Ong, Editor, London

Jaeckel, P, (2005) "Splitting the Core", Working Paper, ABN AMRO.

Moody's KMV, (2004), Early Warning provided by the Moody's KMV expected default frequency credit measure, professional services client report, July 2004

Vasicek O, (1991), *Limiting Loan Loss Probability Distribution*, KMV corporation