J Finan Serv Res (2007) 32:17–38
DOI 10.1007/s10693-007-0011-6

Basel II: Correlation Related Issues

Sanjiv R. Das

Received: 8 October 2006 / Revised: 9 February 2007 / Accepted: 15 May 2007 /
Published online: 28 July 2007

Abstract Basel II aims to aggressively improve on Basel I, and is projected to capitalize on the technological advancements that have permeated the financial industry since Basel I. This paper examines the correlation issues that arise, and provides recommendations on implementation as we move forward. We provide the following results: (1) We demonstrate that fixing asset value correlations by regulators without a specification of business unit granularity and aggregation impacts franchise risk. (2) Loss distributions for credit risk are more sensitive to correlation assumptions that those for market risk; arbitrary, inaccurate correlation specifications can cause large errors in capital requirements. (3) Current regulations do not recognize that credit losses depend on four distinct correlations, not just one. (4) Recovery rates may be determined uniformly across banks. (5) Tail risk comes from LGD correlations and non-Gaussian risks. (6) The 1-year VaR horizon causes distortions especially when regimes and pro-cyclicality are involved. (7) We recommend a quantitative measure for implementing market discipline, the third pillar of the Basel II accord. Therefore, this paper highlights many issues that may be addressed using the tools banks already employ for internal risk management.

Keywords Basel · risk management · correlation

These comments are based on a discussion of various areas of Basel II rules and the US notice of proposed rule-making (NPR). These comments serve the purpose of linking Basel II issues to other published work, of the author and others. The contents of this paper were delivered at the 6th Annual FDIC Conference on Banking in Washington, DC in September 2006. The contents also contain analyses that are original and distinct from other work as well.

S. R. Das (✉)
Santa Clara University, CA 95053, USA
e-mail: srdas@scu.edu

1 Introduction

The Devil is in the Details. No, it's in the Tails.

Basel II introduces two enhancements to the framework of Basel I. First, a tighter link between risk (market, credit and operational) and capital, with the introduction of an internal ratings-based (IRB) approach. For market risk, more refined assessments of risk are made based on the sensitivities of risk positions, than on outstanding notionals. For credit risk, the ratings-based approach will be supplanted by one that looks more carefully at counterparties using internal models. The consideration of operational risk also brings in to the analysis a source of unexpected loss that has received less attention so far, but needs addressing given the potential size of losses from this source of risk. Second, an important conceptual underpinning of Basel II lies in an explicit consideration of correlations in the total risk of portfolios, and consequently, the influence of correlated risk on capital requirements is better exemplified. In this paper, we consider various technical issues in the Basel II framework as they pertain to both, the first and second enhancements, with greater emphasis on the latter. An excellent wide-ranging discussion on Basel II is provided in Gup (2004).

Our approach in this paper is to discuss various issues related to the new Basel II requirements, as well as the NPR (notice of proposed rule making) in the light of current research findings. We first summarize some of the regulations that we will refer to in the course of this comment, and then proceed to present evidence from extant research that suggests more analysis is required of some of the implementation details of Basel II. This comment is agnostic about whether Basel II is likely to result in over or under-capitalization of major financial institutions (quantitative studies provide a prognosis of lower capital requirements). Rather, the attempt is to focus on areas that need further clarification and analysis, with a view to improving the accuracy of the amount of capital maintained.

The change from the old system under Basel I to a new system under Basel II is well-summarized by the comments of Fed Chairman Ben Bernanke at the Federal Reserve Bank of Chicago's 42nd Annual Conference on Bank Structure and Competition, Chicago, Illinois on May 18, 2006. To quote:

> "...the relatively crude method of assigning risk weights to assets, as well as an emphasis on balance-sheet risks as opposed to other risks facing financial firms, limits the overall responsiveness of capital requirements to risk under Basel I, which renders that system increasingly inadequate for supervising the largest and most complex banking organizations. For these organizations, we need to move beyond Basel I to a more risk-sensitive and more comprehensive framework for assessing capital adequacy. Basel II represents the concerted efforts of the supervisory community, in consultation with banks and other stake-holders, to develop such a framework."

There are many issues surrounding the impending implementation of the Basel II IRB (internal ratings based) approach. Proponents for this highlight many advantages such as (a) a reduction in the amount of capital being held, (b) more dynamic and realistic capital adequacy computation, (c) risk-based pricing of products, (d) a means to instill best practices, (e) the introduction of much needed analytical methods, (f) reduction in expected future charge-offs from current feedback into internal risk management systems and rationalization of current levels of risk, (g) reduction in operating expenses from workflow rationalizations arising from the assessment of operating risk (see Buehler et al. 2004), (h) reduction in operating losses, (i) better capital allocation amongst business units within a financial institution, (j) improved corporate governance, and (k) overall lower systemic risk in the financial system.

On the other hand, opponents of the new accord suggest many disadvantages such as (a) a high cost of implementation, (b) competitive disadvantages between banks that are not required to comply and those that have to, (c) competitive imbalances across countries as different national supervisors impose varied levels of compliance, (d) strong opposition to operational risk charges as being a deadweight cost for imposing governance that is already legally mandated, (e) inability to obtain consistent implementation across all institutions, resulting in more noise than accurate determination of risk, (f) the propensity to increase systemic risk if the rules impose distortionary portfolio changes in one same direction across all financial institutions.

To summarize, Basel II envisages three *pillars*: (1) capital adequacy, (2) regulatory review, and (3) market discipline, over three categories of risk, (a) market risk, (b) credit risk and (c) operational risk. In this comment, we will focus mainly on capital adequacy and more specifically, on correlation issues for the market and credit risk components of capital requirements, though we will also comment on technical tools for market discipline.

The previous approach to capital adequacy relied on taking a portfolio of a given size and ascribing to it a risk factor, based on which the capital requirement was imposed from a table. Clearly, this suffers from the deficiency that it ignores portfolio specific risk, that portfolios tend to be quite different in their individual characteristics, even when they are of the same asset class, leverage and maturity. By suggesting that we move on to a value-at-risk (VaR) like system, where the loss distribution is explicitly modeled is clearly going to determine capital adequacy better, provided that the calculations involved and the modeling assumptions are practical and reasonably accurate. However, moving to the IRB approach allows banks greater flexibility in making a wide range of assumptions to "cook" the numbers to achieve internal target capital levels. Yet, one may be optimistic that this is unlikely to occur (a) with more oversight, (b) the fact that the IRB approach recognizes that banks have already been using risk-based capital for almost two decades now, and (c) that this new approach is much more consistent with internal risk management. By all counts, this will reduce the costs of internal and regulatory risk management in the long run, though in the short run, the need to produce both Basel I and II reports is no doubt an onerous imposition.

To summarize, Basel envisions two types of losses: (a) *expected loss* (EL), and (b) *unexpected loss* (UL). If the horizon for the analysis is denoted T, and the current value of a portfolio today (at time t) is $P(t)$, then expected loss (EL) is:

$$\mathrm{EL}(T-t) = E[P(T) - P(t) | P(T) - P(t) < 0]$$

where $P(T)$ is the value of the portfolio/asset at time T. (The asset may have a maturity beyond T). For market risk, the 1996 Market Risk Amendment specifies the horizon $(T-t)$ to be 2 weeks. For credit risk, BCBS (2005) provides a maturity adjustment.

VaR at a level of α (say 1%), is defined as the tail cut off $[P_\alpha(T) - P(t)]$ for which losses in excess of this value will occur with α probability. We write this loss value as $\mathrm{VaR}(\alpha, T-t)$. Unexpected losses are then defined as:

$$\mathrm{UL}(\alpha, T-t) = \mathrm{VaR}(\alpha, T-t) - \mathrm{EL}(T-t)$$

Losses in excess of $\mathrm{VaR}(\alpha, T-t)$ are denoted as *extreme losses* and may also be reserved for. However, the guidelines focus on EL and UL.

Inextricably tied up with the concepts of expected and unexpected loss are the notions of *regulatory* capital and *economic* capital. Loss reserves are meant to buttress expected losses, and economic capital is for unexpected losses. One would also expect that economic capital plays a bigger role in maintaining the credit rating of a financial institution. Again, for intuitive reasons, EL is not sensitive to the shape of the loss distribution as much, whereas UL clearly is. UL is also susceptible to many ills that risk measures like VaR suffer from, such as failure to be a "coherent" risk measure, per Artzner et al. (1999).

This may be a good point at which to recap that whereas VaR has widespread use, it has some well-recognized flaws: (a) It is not a "coherent" risk measure, in that it fails the "sub-additivity" criterion, which simply put, says that a risk measure should always be lower when a portfolio is diversified. In the case of VaR, this is not guaranteed; indeed, taking a weighted average of two portfolios may result in an increase in the risk measure. Intuitively, this occurs because VaR is a percentile measure, and not a moment of the loss distribution. However, in all fairness, as we increase the VaR cut off (i.e. extend further into the tails of the loss distribution), the failure of sub-additivity is much less likely to occur. (b) VaR is very hard to measure because it depends wholly on the tail of the loss distribution. At tail cut offs of 99.99%, it is hard to be confident of its value. There is really no data to validate the efficacy of chosen cut-offs. This is popularly known as the "Star-Trek" problem, i.e., how do we estimate properties of a region of data where we have never even gone before. See Lo (2000) for a nice exposition on this problem. Chorafas (2004), page (xxii) cites a study by Citigroup that claims that sample size impacts the accuracy of VaR, which suggests that the NPR may issue guidelines on sample sizes needed to achieve acceptable estimates of VaR. (c) VaR is often computed under the assumption of normal multivariate distributions of the assets in the portfolio. It is well known that this is an assumption made purely for analytical and computational convenience. Experiments undertaken showed that VaR was 50% higher when the normal distribution was replaced by a student-T distribution with 5 degrees of freedom. Such numbers hint at how easy it is to be undercapitalized. (d) Lo (2000) also points out that VaR is usually based on an unconditional distribution of portfolio P&L, which is satisfactory for passive portfolios but not for actively managed ones. What is needed for active risk management is a conditional measure, incorporating conditional correlations (see Engle 2002).

To get a sense of the magnitudes of *market* risk in the financial sector, it is interesting to examine the data on VaR reported by Jeffery and Chen (2006). They show that the average VaR in the world's 25 leading financial institutions for 2005 was $51.9 million (one-day VaR, at a 99% level). Of this, the biggest component is interest rate VaR, then equity and commodity VaR respectively.

When speaking specifically about *credit losses*, the accord envisages four *risk components*, i.e., (a) probability of default (PD), (b) loss given default (LGD), (c) exposure at default (EAD), and (d) a maturity adjustment (M). For credit losses,

$$\mathrm{EL} = \mathrm{PD} \times \mathrm{LGD} \times \mathrm{EAD} \times f(\mathrm{M}) \tag{1}$$

Formulas are provided for the derivation of risk-weighted assets which depend on estimates of PD, LGD and EAD as well as effective maturity (M; details are provided in BCBS (2005), see Part 2 [*The first pillar: minimum capital requirements*], Section III ["Credit risk—internal ratings based approach"], Section C).

Whereas the document presents the formula as above (more or less), what is hidden is that all these four risk components above may be stochastic and drawn from distributions as well (with correlation amongst them). The formula above suggests that the expected values of these risk components be used to determine EL, with the concomitant result of running afoul of Jensen's inequality, though it is unclear in which direction. However, since it is widely accepted that PD and LGD are positively correlated the impact of Jensen's inequality is most likely adverse. In order to ascertain UL, a distribution of losses needs to be generated under many scenarios accounting for the fact that these inputs vary, and that the occurrence of default is also subject to the specific realization of the value of PD. The actual LGD also may be variable. The devil lies very much in the details here. As we will soon see, credit loss distributions are far more tail dependent than that for market risk, making correlation assumptions difficult to stipulate, validate and implement. Arguing even more theoretically, Jarrow (2006) suggests that Basel II needs to be much more careful in its assessment of system wide debilities, and that unless care is taken, eliminating current capital rules is likely to be fraught with risk.

The Basel II framework suggests two levels of IRB implementation: (1) *foundation*, or F-IRB and (2) *advanced*, or A-IRB. In the former, banks use their own PD, but take LGD and EAD as provided by regulators. In the latter, banks use their internal estimates of all input parameters. As we will see, some of the complication lies in the risk model, and much of it in the input assumptions. In the following sections, we examine various issues that need careful consideration by financial institutions implementing Basel II. The discussion will also relate these technical issues to known empirical realities from the recent literature, so as to develop an understanding of the areas in which Basel II rules offer good risk-based capital assessment, and those in which they do not.

2 Aggregation Level for Business Units

Fixing asset value correlations between business segments based on empirical correlation studies may result in perverse results for the overall capital to be maintained across a franchise. Correlation assumptions need to depend on the choice of aggregation level (granularity) chosen when composing business units, which can make a substantive difference to the computed risk measure at the top level for the institution as a whole. There are two types of granularity that may result in imperfect aggregation of risk: *sector* concentration, and *name* concentration. We focus more on the former.

Kuritzkes et al. (2003) suggest that instead of a silo approach to risk measurement (where capital adequacy is imposed at the business unit level with no diversification), capital requirements computed at the level of an average banking and insurance conglomerate would be 5–10% lower than the sum of capital for each unit. They suggest that a risk factor approach be used instead of a business unit approach in compartmentalizing risk.

Aggregation and granularity issues also arise in the the realm of name concentration, when a bank has multiple exposures to the same obligor. These issues are addressed in a series of papers by Gordy (2003, 2004), and Gordy and Lütkebohmert (2006). These papers also examine when portfolio invariance matters. Incremental risk of an additional position depends of course on the portfolio to which it is being added, and hence, different business unit compositions can result in different measurements of marginal risk contributions and different calculations of additional capital required.

In contrast, we focus on a specific aspect of these issues. We are interested in how different choices of business units (totaling up to the same overall franchise) might result in different levels of capital requirement. Hence, we look specifically as how total franchise risk can depend on the breakdown into business units. If the entire franchise is atomized into individual risk positions, one for each transaction on the books, we will capture all the risk as long as we consider the covariance matrix of all these positions. The way in which aggregation is undertaken results in different outcomes because it results in using specific subsets of this overall covariance matrix. Different subsets results in different aggregate risk totals. We show that the risk capital required will be a u-shaped function of the granularity of the franchise. Therefore, complementing some of the previous research (for example, BCBS 2004) that looks at inadequate diversification from sector concentration, resulting in higher capital requirements, we analyze an already diversified firm, and show that, ceteris paribus, differential allocation to alternate groupings of business units can result in differences in total capital required.

In a single risk factor framework, each transaction has systematic and idiosyncratic risk. At one extreme, we may have each transaction as a separate portfolio or business unit (perfect disaggregation). As transactions are grouped into sub-portfolios (units), diversification reduces the risk within each sub-portfolio. Overall risk for the franchise should remain unchanged, since sub-portfolios now become more correlated as the ratio of systematic risk to idiosyncratic risk across sub-portfolios will increase. This implies that the correlation between asset classes (sub-portfolios) depends on the extent of granularity chosen. In a situation where the correlations are *fixed* based on empirical estimates independent of granularity, aggregate franchise risk is in fact distorted. The obvious reason for this is that if the regulator computes empirical correlations between asset classes based on empirical evidence for a *given* level of granularity, banks may then choose a lower level of granularity for their portfolios, and thus in fact keep less capital than they should. A simple numerical experiment here shows that there is an optimal level of granularity that a franchise may choose to minimize its required capital. Clearly, this concern has been recognized within the NPR, since it suggests as high a level of granularity as possible.

To illustrate the critical nature of the granularity and aggregation decision in determining capital adequacy, we conduct the following experiment. Assume we have $n = 2^{10} = 1{,}024$ assets in our portfolio and each asset has mean value 0 and variance 1, i.e. we may describe them as standard normal variables (this is without loss of generality). We also assume that the correlation between these variables is the same for each pair, and is denoted ρ. Hence the covariance matrix of asset values (Σ) is of dimension $n \times n$ with the value 1 on the diagonal and ρ off-diagonal. Assume an equally weighted portfolio (P) of these assets, i.e., w is a vector of weights, each of value $1/n$. The mean value of this portfolio is 0 and its variance is $\sigma^2 = w' \ \Sigma \ w$. We will compute the EL, UL, and VaR of P.

$$\mathrm{EL} = \int_{-\infty}^{0} P \frac{1}{\sqrt{2\pi\sigma^2}} \exp(-P^2/2) \ dP$$

where the formula above is the expected value of P conditional on it being less than 0, and assuming it is distributed normally.

The 1% VaR of this portfolio is obtained by inverting the cumulative normal distribution for the left tail area of 0.01. Finally, the UL is determined using EL and VaR. If we assume that $\rho = 0.5$, we obtain the following values:

$$\mathrm{EL} = 0.2822, \quad \mathrm{UL} = 1.3635, \quad \mathrm{VaR} = 1.6458$$

where of course, the values are taken with positive sign (absolute value) since we are interested in the loss distribution, even though the integral above results in a negative value. We might imagine that what we have here are 1,024 separate business portfolios and that we aggregate them all equally weighted into one enterprise portfolio and compute the risk measures above. This set up assumes a very high level of granularity, i.e. each asset is a distinct business unit.

Next, suppose we construct each business unit as comprising $m = 2$ assets, so that the 1,024 securities are devolved into $n = 512$ business units or portfolios. Our basis for computing EL, UL and VaR now requires the covariance matrix of these 512 portfolios. Note that each portfolio has mean 0 as before, but the variance is $y' \ \Omega \ y$ where y is a vector of length m of values $1/m$. Ω is a matrix of dimension $m \times m$, with 1s on the diagonal and ρ off-diagonal.

We construct the covariance matrix Σ of the entire franchise (this is of dimension $n \times n$ or 512×512) [the correlation parameter ρ here between portfolios may actually vary, depending on the actual factor model used. But our goal here is more to demonstrate the perverseness of risk measurement for varying aggregation levels]. The diagonal elements are $y' \ \Omega \ y$, and the off-diagonal elements are $\rho(y' \ \Omega \ y)$. This is the crux of the problem at hand, the parameter ρ is not adjusted for the change in granularity. Note that this will be less than ρ since the pairwise correlation between portfolios of assets will be lower than the pairwise correlation of the assets themselves. For the same values this correlation is now 0.375, not 0.5 as for individual assets. Hence, the level of granularity here is halved, each business unit is two assets. Also, the aggregation is based on a different correlation matrix. After running the computations as before, we obtain the following values:

$$\mathrm{EL} = 0.2445, \quad \mathrm{UL} = 1.1814, \quad \mathrm{VaR} = 1.4260$$

The capital required here is lower, as all three risk measures are smaller than before. This is not surprising as the correlation across business units has been dampened as they are made up of portfolios. We undertook the same computation for a changing number of portfolios of the 1,024 assets, by dividing the number of business units progressively by 2, and multiplying the number of assets within each portfolio by 2. The results for EL, UL, and VaR are shown in Table 1.

The results are interesting. As we reduce the level of granularity the risk measures fall. This is because there are two types of diversification: (a) diversification *within* portfolio or business unit, and (b) diversification *across* units. When the portfolio is the same as the asset, there is no within unit diversification, only across units. As we increase the number of assets in each portfolio, we get diversification within unit, and also across unit. Think of the original covariance matrix being halved in dimension and block diagonalized so as to lose some of the correlation between individual securities *across* units. Hence, this leads to *Problem #1*, i.e. that the level of granularity affects measures of risk, even though the total risk has not been changed.

As granularity is reduced further, again there is a trade-off between diversification within and across portfolios, until at some point, we begin to lose diversification across units, as the number of portfolios becomes too small, and the risk measures begin to rise once again. We can see that there is a material difference between the risk measures at varied granularity levels. This leads to *Problem #2*, i.e. when the regulators (internal or external) provide asset value correlation (AVC) levels, what granularity level do they have in mind?

If we assume that the correct level of the risk measures is based on the highest level of granularity, then is there some way in which we can undertake a mathematical fix (or hack) to bring the values back to the risk measures as based on the highest level of granularity? It turns out that in the case of our example, this is easy to do. All that we need is to reset the calculations assuming that the correlations between business units is replaced by the correlation levels between individual assets (that is, in the third column of Table 1, set Corr=0.5 for all rows). This is clearly wrong but it does counteract the problem! Of course, this is easily done in our example, because we have assumed all assets have the same distribution and the same pairwise correlations. This raises *Problem #3*, i.e. how do we provide correct AVC correlations

Table 1 Expected loss, unexpected loss and value-at-risk for varying levels of granularity and aggregation. The first column shows the number of business units, and the second one the number of assets within each unit. Each asset has a standard normal distribution. "Corr" is the average pairwise correlation between portfolio values

# Portfolios	# Within portfolio	Corr	EL	UL	VaR
1024	1	0.5000	0.2822	1.3635	1.6458
512	2	0.3750	0.2445	1.1814	1.4260
256	4	0.3125	0.2235	1.0796	1.3030
128	8	0.2812	0.2124	1.0261	1.2385
64	16	0.2656	0.2072	1.0011	1.2083
32	32	0.2578	0.2057	0.9938	1.1995
16	64	0.2539	0.2072	1.0011	1.2083
8	128	0.2520	0.2124	1.0261	1.2385
4	256	0.2510	0.2235	1.0796	1.3030
2	512	0.2505	0.2445	1.1814	1.4260
1	1024	0.5000	0.2822	1.3635	1.6458

when the assets are highly heterogeneous. These are already being stated, but as our example above shows, there is bound to be an ad-hoc component to it.

The simple correction required is to increase ρ across business units as the level of granularity declines. But by how much? In our example, where all assets have the same distribution, it is easy to compute the correction. But when the assets within each unit are heterogeneous, there is no simple way to do this. Again, the NPR clearly envisages this problem in requiring that units be defined for highly homogeneous assets. In short, granularity of risk measures complicates risk aggregation on account of correlation assumptions. As the number of assets in the VaR simulation becomes very large, several useful asymptotic approaches may be applied (see Gordy 2003, 2004; Gordy and Lütkebohmert 2006).

3 Correlation Sensitivity of Credit Portfolios

For credit portfolios, risk measures, based on loss distributions are highly sensitive to the correlation parameter. Credit portfolios are essentially based on binary outcomes, and hence the joint distribution is quite different than with portfolios where the outcomes reside on a wide range of values. Intuitively we will see that a portfolio of binary outcomes has a distribution that changes very quickly when correlations change than say, a portfolio where the assets are distributed multivariate normal.

In the previous section, we saw that the risk measures EL, UL and VaR are sensitive to aggregation level. The analysis there was simple and assumed distributions over a continuous range of values. However, when dealing with credit losses, the value tends to be either zero (no loss) or a loss value within some range. Intuitively, each asset follows a Bernoulli distribution, with one outcome being zero. When we construct portfolios of such assets, the distributions become even more sensitive to correlation assumptions, implying that the risk measures will also be much more variable when correlations are changed.

To illustrate, we work in the standard asymptotic single risk factor (ASRF) framework that is now very popular in an analyzing correlated default risk. Assume there are n assets in a portfolio. Each asset is identical with a Bernoulli outcome over values $\{0, \text{LGD}\}$ with probability $\{1-\text{PD}, \text{PD}\}$ respectively. For normalization assume that $\text{EAD}=1$, and that $\text{LGD}=1$. In this example, we assume that the PD is stochastic and that all other input variables are constant.

In order to inject correlation amongst defaults, we examine the following set up. Assume that the n assets each have an underlying value process as follows:

$$x_i = \sqrt{\rho}\, z + \sqrt{1-\rho}\, e_i, \quad z, e_i \sim N(0,1), \quad \forall i.$$

Hence, $E(x_i)=0$, and $\text{Var}(x_i)=1$, for all assets, assuming that z is independent of all e_i, and that the e_is are independent. Here, the correlation ($\rho = \text{Cov}[x_i, x_j]$) amongst the assets is generated from the common random variable z. Note that $\text{Cov}(e_i, e_j)=0$ for all pairs (i, j). Since means are zero and variances are 1, the covariance is also the correlation. This is a standard set up and for more examples of credit loss computations, see Kupiec (2005).

For asset i, default occurs if $N(x_i) < \text{PD}$, where $N(.)$ stands for the cumulative normal distribution. As is well known, it is easier to build up the loss distribution if we condition on various values of z. Suppose we fix a value of z. Then the probability of default, conditional on z is denoted $PD|z$, and is

$$\begin{aligned}
\text{PD} \mid z &= \text{Prob}[N(x_i) < \text{PD} \mid z] \\
&= \text{Prob}[x_i < N^{-1}(\text{PD}) \mid z] \\
&= \text{Prob}\left[\sqrt{\rho}\, z + \sqrt{1-\rho}\, e_i < N^{-1}(\text{PD})|z\right] \\
&= \text{Prob}\left[e_i < \frac{N^{-1}(\text{PD}) - \sqrt{\rho}\, z}{\sqrt{1-\rho}}|z\right] \\
&= N\left[\frac{N^{-1}(\text{PD}) - \sqrt{\rho}\, z}{\sqrt{1-\rho}}|z\right] \\
&\equiv q_z
\end{aligned}$$

The probability that there are m losses from n firms, conditional on z is denoted $p_z(m)$, given by the binomial formula

$$p_z(m) = \binom{n}{m} q_z^m (1-q_z)^{n-m}, \quad m = 0...n.$$

Noting that $z \sim N(0, 1)$ we can integrate it out to get the full loss distribution, with the probability of m losses:

$$p(m) = \int_{-\infty}^{\infty} p_z(m)\, \phi(z)\, dz, \quad m = 0...n.$$

Table 2 Risk measures for varying default correlation. The PD for each firm is 5% and the number of identical firms is 100. The expected loss should be exactly 5.00 for all correlation levels, and the tiny discrepancy comes from numerical rounding error. The last column contains the adjustment term from the formula on page 405 of the draft NPR, i.e. $N\left[\frac{N^{-1}(\text{PD}) - \sqrt{\rho}\, N^{-1}(0.999)}{\sqrt{1-\rho}}\right]$. We can see how it varies with correlation

Corr	EL	UL	CVar	Kadj
0.00	5.0000	5.2046	10.2046	0.0500
0.10	4.9991	13.1910	18.1902	0.2408
0.20	4.9984	20.7485	25.7469	0.3844
0.21	4.9984	21.5080	26.5064	0.3985
0.22	4.9983	22.2602	27.2585	0.4124
0.23	4.9982	23.0061	28.0044	0.4264
0.24	4.9982	23.7793	28.7775	0.4403
0.25	4.9981	24.5474	29.5455	0.4542
0.26	4.9981	25.3110	30.3090	0.4680
0.27	4.9980	26.0713	31.0693	0.4817
0.28	4.9980	26.8482	31.8461	0.4955
0.29	4.9979	27.6300	32.6279	0.5091
0.30	4.9979	28.4099	33.4078	0.5227
0.40	4.9975	36.4720	41.4695	0.6553
0.50	4.9973	45.1927	50.1900	0.7776
0.60	4.9972	54.8697	59.8670	0.8818

where $\phi(t)$ is the normal pdf. This is easily computed using a fast quadrature routine or discrete integral. Once we have the loss distribution we can compute EL, UL and CVaR (credit VaR). Note that this approach is fairly standard and correctly produces credit loss distributions with the desired correlation. The risk measures are shown in Table 2. It is evident that the UL risk measure (and hence economic capital) is very sensitive to correlation assumptions. To get a visual feel for how quickly the loss distributions change, see Fig. 1. Also note the last column in Table 2. It contains the term that varies as correlation changes in the capital formula from page 405 of the draft NPR. It complements this analysis in that the correlation adjustment tracks the vastly changing risk measures quite well.

It is therefore apparent that portfolio risk with assets that have default risk is very sensitive to credit correlations, far more than is the case with correlations pertaining to market risk, such as in equity portfolios. Specifying correlations correctly is paramount because even small changes in correlation result in drastic changes in the loss distribution as shown in Fig 1. From a regulator's point of view, it is therefore important that banks be encouraged to invest effort in correctly assessing credit correlations, and that these assumptions be carefully reviewed.

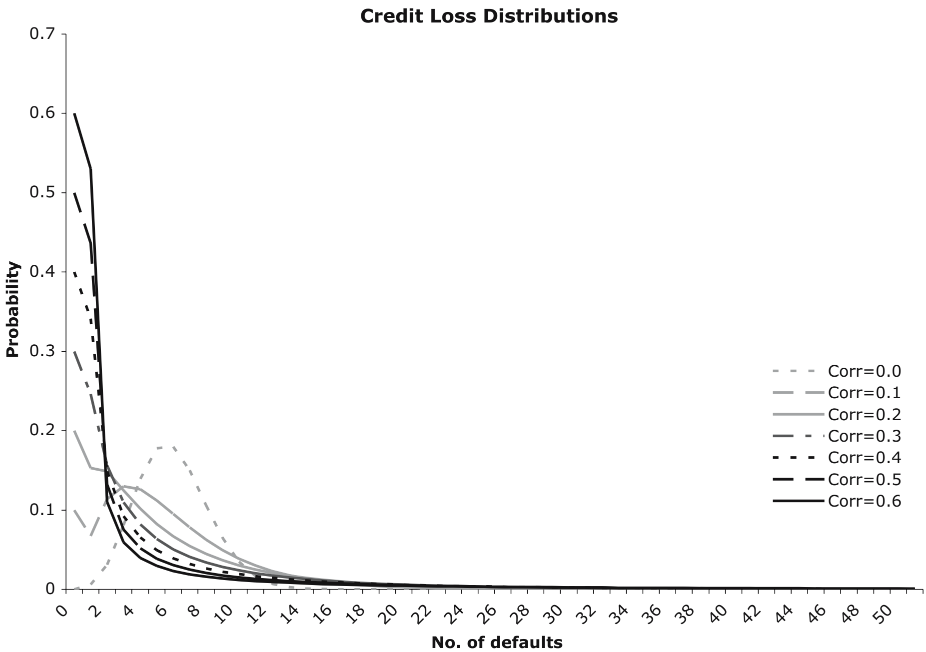

Figure 1 Credit loss distributions under varied default correlation levels. We only present the loss levels out to 50 on the x-axis, even though the maximum number of defaults is 100, as the probabilities become very low thereafter. Note the distribution is mostly symmetric under the zero correlation assumption, and then becomes sharply skewed rapidly as we increase the level of correlation. At the left most edge of the graph, the line for zero correlation is the lowest, and the one for a correlation of 0.6 is the highest

4 Asset Value Correlations (AVC)

The agencies involved in the formulation of the NPR have mandated two principles that need some discussion. First, the concept of *portfolio invariance*, and second, the use of *correlation factors*. The first allows each individual position's risk exposure to be calculated without accounting for other risks that might be taken in the remainder of the franchise. The second adjusts capital for the additional risk that arises from correlations. The *first question* that arises is whether the correlation adjustment that is made accounts properly for the sensitivity of the risk measures to correlation changes. As we have seen in Table 2, this certainly appears to be so.

An important *second question* is whether the assumption that default probabilities in low-PD portfolios are more correlated than in high-PD portfolios is a valid one [this question is raised for discussion on page 67 of the draft NPR]. There is evidence that supports this assumption, which emphasizes that if we examine PDs, then since bigger, safer firms have more relative systematic risk than smaller, riskier firms, their PDs tend to move together to a greater extent. However, the analysis appears to be incomplete to some extent, because it looks only at the co-movement in PDs, but not also at the event of default conditional on the PDs.

It is useful here to take the viewpoint of doubly-stochastic reduced-form models. In these models, default depends on two stochastic processes, (1) one process driving default probabilities, and (2) conditional on a PD, another random variable driving the event of default. Correlated default occurs on account of either or both of these stochastic processes being correlated across firms. In short, defaults are correlated because firms' PDs are correlated. Defaults may also be correlated even when PDs are independent, if contagion effects exist, and the default of one firm triggers the default of others.

Low-PD firms tend to display higher PD correlations than high-PD firms. This is evidenced in a study by Das et al. (2006). Figure 2 contains a graphical reproduction of Table 4 from this paper. It shows that high quality firms have higher PD correlation than low quality firms, across four economic regimes. This confirms and supports the ideas embedded in the NPR regarding adjusting correlation for economic regimes.

Further, we need to consider whether defaults might be correlated differently for the second part of the doubly stochastic reduced form model. In other words, are contagion effects more prevalent amongst high-PD firms as opposed to low-PD ones? The presence of contagion (or frailty) effects has been empirically confirmed in Das et al. (2007). Whether these are more prevalent amongst high or low quality firms is an open question that requires further empirical analysis. From a historical perspective, the late 1980s were a time when contagion might arguably have been prevalent amongst high-PD firms. But in the early 2000s, major contagion effects were evidenced amongst fairly large, well established firms.

In the context of contagion, we see that in periods when PDs were high, as in Fig. 2, (the first and fourth periods), default probability correlations tend to be 2 to 4 times as high as in periods with low PDs. Thus, in downturn scenarios, UL may be based on a correlation level of 0.40 versus normal times, with average correlations of 0.10. From Table 2 we see that UL will change by a factor of 2.7, which would also be the change in economic capital required. Hence, it seems appropriate to allow capital to be dynamically adjusted in periods of economic downturn, rather than

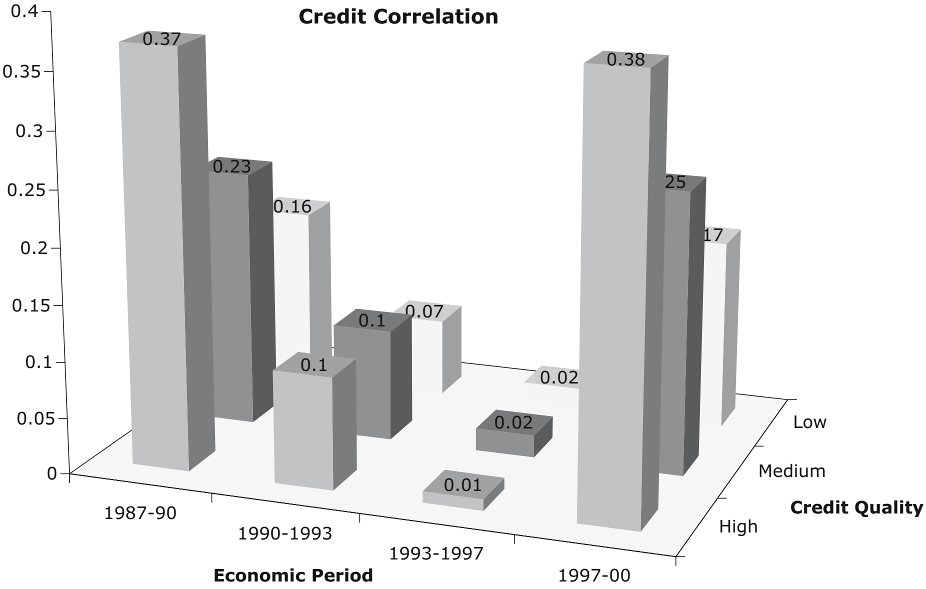

Figure 2 Default probability correlations. Based on the data of Table 4 from Das et al. (2006) showing that low-PD (high quality) firms have higher default probability correlations than high-PD (low quality) firms. We also see that in two of the four economic regimes, i.e. the post-Millken period (1987–1990) and the current dot-com bust period (1997–2000), not only were default levels high, but default probability correlations are also higher than in the other two economic periods

build this into capital requirements on a continuous basis as envisaged in Table 2, page 405 of the NPR.

This naturally raises a *third* question of how to detect a down cycle, which is characterized in industry circles as the "detection of pro-cyclicality." We submit here that this is not as infeasible as it might have been a few years ago. We now have evidence that models for aggregate default intensity correlate strongly with actual default levels. Give this, we might be able to use aggregate PD measures to assess when a regime shift has occurred. As is shown in Duffie et al. (2007) and in Das et al. (2007) (see Fig. 2 there), the model for default intensity does a good job of tracking defaults. Further, Das et al. (2006) provide a regime-switching model for default probability, in which transition probabilities into a high-default regime are estimated. Hence, we may be able to use various models such as the one by Duffie et al. (2007) for detecting pro-cyclicality.

A *fourth* question that arises is whether the specific assumption of low (high)-PD corresponding to high (low)-credit correlation will distort the amount of economic capital required. Noting from Table 2 that UL (economic capital requirement) is very sensitive to correlation might well result in a high-PD but low correlation portfolio requiring less economic capital than a low-PD high correlation portfolio, even after the higher risk capital set against high-PD assets. Again, this highlights that correlation assumptions may be tricky, resulting in non-intuitive capital requirements.

5 Loss Given Default

The determination of LGD required for the EL computation is a difficult issue. See Schuermann (2004) for a comprehensive guide to various issues pertaining to LGD as related to Basel II. We are in need of models that allow us to determine a forecast of recovery conditional on default. One such model that makes use of easily available information at a given point of time has been developed by Das and Hanouna (2006). The model is flexible and uses the information in CDS spreads to determine both default probability and recovery. The model may also be applied on average to sector spreads if need be to obtain a coarser estimate of recovery rates that may be more amenable to regulatory use. There are several other papers that provide insights into recovery rates, that may also be used to guide regulators in developing methodology for recovery rates that will be input into IRB models. Altman et al. (2005) found considerable variability in a time series of default rates and recovery levels in the US corporate bond market. They document a statistically significant negative relation between these two variables. Whereas this analysis relates to the real-world probability measure, pricing would require the risk-neutral measure, and Unal et al. (2001) develop a simple model to infer these from various securities in the firm, accounting also for APR violations.[1]

In a reduced-form default model, we may use CDS spreads to extract the term structure of forward default probabilities (λ), making a simplifying assumption about recovery rates (ϕ, or LGD $= 1 - \phi$). Suppose we are given the term structure of CDS spreads on any day. If we fix recovery rates to be a constant ϕ, then simple bootstrapping allows us to determine the term structure of default probabilities. However, we need to know ϕ a-priori, requiring us to make exogenous assumptions about its value. Indeed, Basel II suggests that regulators provide this value to financial institutions in some implementation versions.

Das and Hanouna (2006) provide an alternate way to determine ϕ using additional data from the equity markets. The brief outline of the algorithm is as follows. First, fit the parameters of the Merton (1974) model to the values of the stock price and stock volatility. This provides the firm value and firm volatility. Second, use the Merton model to express the recovery rate $\phi(T)$ for each maturity T as a function of the default probability $\lambda(T)$; this expression is as follows:

$$\phi(T) = e^{rT} \frac{V}{F} \frac{N[-d_1]}{\lambda(T)}, \quad d_1 = [\ln(V/F) + (r + 0.5\sigma^2)T]/[\sigma\sqrt{T}]$$

where r is the risk free rate, V is the value of the firm, F the face value of debt, $N[.]$ is the cumulative normal distribution function, and σ is firm volatility. Third, use this expression for $\phi(T)$ in the bootstrapping procedure to determine the values of $\lambda(T)$ for all T.

Firms that choose to adopt the A-IRB approach may use this framework to determine the recovery rates for individual names in their portfolios. Regulators

[1] Many other papers undertake similar work. Zhang (2003) identifies default intensities and recovery rates in a reduced form model, applied to Argentine sovereign debt; recovery rates are estimated to be approximately 25%, the common number used in the market. Pan and Singleton (2005), using a panel of sovereign spreads on three countries (Mexico, Russia and Turkey) to identify recovery rates and default intensities assuming recovery of face value.

may also use the approach to determine recovery rates for dissemination under the F-IRB approach. Adoption of a simple model such as this may therefore enable a standardized approach to LGD across financial institutions.

6 Contagious Interaction of PD and LGD

Complementing the findings of Altman et al. (2005), the Das and Hanouna (2006) algorithm imposes a negative correlation between PDs and recovery rates, and then estimates the extent of this negative relationship. The algorithm is applied to 3,130 firms over the period 2000–2002, and λ and ϕ are aggregated for all firms (equally-weighted); an inverse relationship between aggregate default probabilities and recoveries is noted. The correlation of default probability and recovery in the cross-section of firms within each month evidences levels of negative correlation of -0.3 to -0.6.

From 2000 to 2002, as default rates rose, recovery rates fell, *and* the correlation between the two became more negative. The relevant implications of these results for Basel II are that (a) the correlation between PD and LGD is important in the application of the $\text{EL} = \text{PD} \times \text{LGD} \times \text{EAD} \times M$ equation. (b) We know from Das et al. (2006) that when overall PD levels rise, their correlations increase. One must take care to distinguish (a) the fact that low-PD firms have higher PD-correlations than high-PD firms from the (b) empirical observation that when economy-wide levels of PDs rise, overall levels of PD-correlation also rise, maintaining the relationship in (a).

From Das et al. (2007) that there is additional contagion correlation detected even after conditioning on PDs. As evidenced in Das and Hanouna (2006), recovery rates become increasingly negatively correlated with PDs as default levels rise, resulting in correlated LGDs (another form of contagion not recognized earlier). Hence, not only do defaults cluster, but when they do, LGDs cluster as well. This has important implications for capital adequacy, and biases capital requirements higher.

Therefore, we now have three different sources of correlation to deal with in the framework of reduced form models. First, there is the correlation between default probabilities of various counterparties. Second, the correlation between PD and LGD is known to be negative, and therefore, results in greater capital requirements. Third, is the contagion effect, where the onset of some defaults triggers more defaults.

And finally, in addition to these sources of correlation within the realm of credit risk, there is the interaction of market risk and credit risk as well. The sign of this correlation tends to be adverse as well. When market risk increases, the three credit correlations are also higher.

7 Non-Gaussian Distributions

Much of the regulatory framework for Basel II relies on a single risk factor Gaussian framework. By using different joint distributions, we may assess the impact of incorrectly adopting the Gaussian model. In Das and Geng (2004) different copulas were applied to the PDs from Moodys over a 14-year period (1987–2000). The industry standard Gaussian copula with normal marginal distributions was found to be inferior to the Clayton copula model with double exponential marginals. Also, the tail of the Clayton copula loss distribution is seen to be much fatter than that of the student *T*. Hence, the Gaussian model will understate the amount of capital required to be maintained. Regulators would be well advised to provide an additional adjustment for non-normality in risk distributions.

8 Does Accounting for Regimes Increase or Decrease Regulatory Capital?

As we have seen, credit risk can vary substantially, both in the level of risk and in credit correlations across economic regimes. Correct maintenance of capital in a regimes-based model comes with complications, as we will now see.

With a time horizon for VaR of 1 year, capital requirements may be mis-stated when regime switching is modeled under the Basel II framework. The problem arises in the case of longish risk horizons (e.g. 1 year), and results in distortions when the portfolio may be modified in time frames less than the VaR horizon. Consider the following thought experiment. Say we are currently in a low risk regime in the economy. Also assume that we can, within a reasonable time (say 1 month) make substantive changes to the portfolio to mitigate risk. Then, accounting for a possible regime shift that might occur in 1 year, where a bad regime is feasible, will result in keeping more capital than is necessary because the probability of switching to the bad regime before the portfolio can be immunized is over-stated. On the other hand, if we are in a high risk regime, where there is a likely switch into a low risk one, will result in keeping less capital than is currently necessary. Therefore, sometimes we keep too much capital and at other times too little. Either way, we always keep incorrect amounts of capital.

Gore (2006) provides a discussion that relates to this issue in the context of retail banking risk. The analysis suggests that it might be best to use horizons appropriate for each business segment, rather than a fixed horizon of 1 year. Segments with low liquidity and longer times to restructure will attract more capital, which also correctly accounts for the liquidity risk in the product line. On the other hand, businesses that engage in liquid transactions, will naturally be more manageable and the liquidity effect will be small, resulting in keeping lesser amounts of capital. A one-horizon fits all approach clearly has its problems, and in particular, complicates keeping correct capital in regime switching environments.

9 Regulatory Safeguards

9.1 Floors on Capital Reductions

The BIS press release of 10th July, 2002, states:

> "More fundamentally, the Committee is proposing to alter the structure of the minimum floor capital requirements in the revised Accord. Under the new approach, there will be a single capital floor for the first two years following implementation of the new Accord. This floor will be based on calculations using the rules of the existing Accord. Beginning year-end 2006 and during

> the first year following implementation, IRB capital requirements for credit risk together with operational risk capital charges cannot fall below 90% of the current minimum required, and in the second year, the minimum will be 80% of this level. Should problems emerge during this period, the Committee will seek to take appropriate measures to address them, and, in particular, will be prepared to keep the floor in place beyond 2008 if necessary."

First, this has implications for the incentives to implement the new IRB based capital requirements, as there is a floor on the benefit that might be attained from moving to the IRB standard. Banks that are likely to have only a small reduction from moving to IRB will find that the benefits from capital requirement reductions might be overwhelmed by the costs of implementing the new Basel II standard.

A *second* effect applies to banks that will experience large reductions in risk capital were they to use the new IRB approach. Such banks will inevitably be disappointed with the floors being placed on capital.

Third, there are many points of tension between the old and new requirements. One might easily imagine circumstances where the risk weights lead to banks that have diversified their portfolios effectively using modern quantitative methods being disappointed when their lower risk levels are not rewarded by an actual reduction in capital required when they hit the floor. This might therefore, disincentivize the introduction of modern risk management methods. The floor requirement also penalizes banks that take active measures to reduce the risk of their franchises, even as they move towards the new IRB approach.

Fourth, it is unclear as to what the guidelines are for the national supervisor to assess the performance of banks so as to release them from the floor capital requirement at the end of the initial 3-year period.

Fifth, banks will keep reserves for meeting EL and also economic capital for further risk. But because we are moving to the IRB approach, the amount of capital to be maintained becomes much more variable given changes in the underlying variables that drive risk even when the portfolio composition does not change. Hence, there is an aspect of the floor requirement that is surely useful, in that it smoothes out fluctuations in capital since a bank already at the floor would not need to keep a reserve buffer given that it was already holding excess capital.

It is clear from the regulator's point of view that the transitional floor requirements are a way of implementing the Basel II framework in a "controlled" environment. Hence, one should not be too critical of the idea. Yet, we do need to take with a pinch of salt the alacrity with which regulators profess they will review their guidelines and remain flexible on changing the norms if they feel that there is a material reduction in capital requirements, failing which the banks would be exposed to unnecessary hardship as a consequence of the transitional floor requirements. Clearly, regulators would like banks to hold more low-risk assets, which did not occur under Basel I guidelines. Given this, the floor capital requirement does not point incentives in this direction.

9.2 Maintenance of Minimum Leverage Requirement

Rules also stipulate a minimum leverage ratio, defined as Tier 1 Capital divided by the adjusted quarterly average Total Assets, after adjustments. The leverage ratio required is a minimum of 3–4% (Tier 1 capital divided by average total consolidated assets. Average total consolidated assets equals quarterly average assets from a bank's most recent Call Report less goodwill and other intangible assets). Banking organizations must maintain a leverage capital ratio of at least five percent to be classified as well-capitalized.

This is over and above the Tier 1 capital ratio of 4% (Tier 1 capital divided by risk-weighted assets) and a Total Capital ratio of 8% (the sum of Tier 1 and Tier 2 capital divided by risk-weighted assets). A well-capitalized institution maintains capital ratios 2% higher than the required guidelines.

The minimum leverage ratio does not account for off balance sheet assets and is likely to become increasingly redundant. One envisages a gradual phase-out of this measure.

How does one include leverage from off balance-sheet positions such as that from derivatives? For example a long position in a call option may be transferred from off balance-sheet to on balance-sheet before computing the leverage ratio. This may be done by recognizing that the option is decomposable into a long position in equity and a short position in a loan. The equity position may then be added to the denominator of the ratio. Such decompositions are non-trivial across a large portfolio but will eventually enable us to establish correctly what leverage representations are especially in the case of an institutional environment in which derivatives are playing an increasing role.

10 A Proposal for Market Discipline

One of the pillars of the new Basel II accord is that of market discipline. A simple approach that may be added to the NPR is that banks also report their "distance-to-default" (DTD) as per the model of Merton (1974). All banks would then be required to maintain a minimum DTD, and if this fell below the acceptable levels, then the banks would need to re-capitalize in order to comply.

Regulatory involvement would require the setting up of this level of DTD. We note that one single level of DTD can apply to all banks, as the DTD is a volatility and leverage adjusted measure, which accommodates differences across banks. Because it is a normalized measure, it is possible to equalize competitive differences across banks and is therefore a possibly useful approach. It is also based on market information and allows risk management of the banking system to be tied to the risk preferences of investors as well.

Regulators may use historical data on defaulted banks to assess the levels of DTD that are "critical" given the target failure rates the FDIC is willing to accept.

Regulators will also need to adjust the DTD limit for economic regime. Given that downturn regimes are characterized by increased default correlations, clearly DTD limits need to be changed so that contagion is not permitted to take root in the banking system. The ample research that now exists on correlated default may now be brought to bear in such a study. Such a measure is more transparent and also consistent with the risk management approaches that banks are more comfortable with.

For example, consider the financials of Citigroup as of December 2006. We examine the data in the SEC filings on the EDGAR system. The book value of equity (in millions of dollars) is 119.8, and short-term and long-term liabilities are 1,476.0

and 288.5 respectively. Using the KMV model, we assume a horizon of 1 year, and also set the effective liabilities to be the sum of the short-term liabilities plus one-half the long-term ones (for a total of 1,620.25). The book value per share is reported at \$24.48, and hence, the number of shares outstanding is 4.95 million. The stock price is \$55.70 and the 1,000-trading day historical volatility (used in the same way as in the `CreditGrades` model) is 17.3%. We applied the inversion technique in the Merton (1974) model to ascertain the asset value per share ($A = \$366.96$) and the asset volatility ($\sigma_A = 2.63\%$). With these in hand, we computed the distance-to-default (DTD), amounting to 4.1439 standard deviations of firm value away from default. This risk neutral probability of default is extremely small, implying that the statistical probability of default is even smaller. Thus, Citigroup clearly need just meet the minimum capital standards. At this DTD, the amount of risk-based capital is also negligible.[2] We might think of this as a "top-down" approach to capital requirements. This is easily reported and also provides another point of comparison with the more detailed "bottom-up" approach. For amplification of this idea, see Merton and Perold (1993).

The application of the Merton (1974) model requires the correct amount of debt on the balance sheet, and in order to assess this, off balance-sheet items will also have to be correctly factored into the analysis. Assessing a bank's liability structure for the computation of the DTD measure may be complicated, yet is a fruitful avenue for further research, especially in the light of recent findings that the most significant variable in default prediction models is distance-to-default (see Duffie et al. 2007). One suggestion would be to compute liabilities using SEC filings as illustrated above, and then increase the liabilities by the amount of the net negative mark-to-market values of off-balance sheet contracts.

Whereas DTD may be an early warning signal, highlighting the need for market discipline, especially if the metric is reflected in the credit default swap spreads, other interesting approaches have also been proposed. Flannery (2005) shows how securities such as reverse convertibles may be used to impose discipline, not as metrics signaling the need for market intervention, but directly through changes in the capital structure.

11 Summary

The discussion in this paper recognizes various technical issues that need careful consideration in the implementation of Basel II, or similar risk-based capital systems.

[2]Merton (1977) showed that risk-based capital per dollar of liabilities for a financial or depository institution was the same as a put option on the bank's assets A with a strike price of the liabilities L plus interest thereon, i.e. Le^{rT}, where T is the maturity of the liabilities. This liability insurance is equal to risk-based capital C.

$$C = N(d_2) - \frac{A}{L} N(d_1)$$

where

$$d_1 = \frac{\ln(L/A) - 0.5\sigma^2 T}{\sigma\sqrt{T}}, \quad d_2 = d_1 + \sigma\sqrt{T}.$$

We focused on the extant literature on risk management, especially as pertains to asset correlations. From a managerial viewpoint, we provide the following concise list of issues that were discussed in the paper.

1. Whereas regulators may provide correlations between asset classes to be used in risk analysis, we have shown that correlations are only meaningful when related to the *granularity* of the portfolios within each asset class, and to the way in which these portfolios or businesses are aggregated into the risk of the entire financial institution. One way to mitigate this problem is to define asset classes more narrowly. A second approach might be to provide an underlying factor structure onto which all assets are projected, thereby allowing factor correlations to drive the joint risk of the entire bank. Since this factor approach is well understood in many markets, such as equity and bond markets, securitizations of credit securities, etc., the Basel II framework may be extended to consider this approach through implementation guidelines in the NPR.
2. We have shown that credit loss distributions are much more sensitive to changes in correlation assumptions than loss distributions for market risk. Hence, care is needed in setting credit correlation parameters in Basel II.
3. Basel II oversimplifies credit correlations, assuming a single form of correlation impacts credit losses. We have shown that in fact there are four different correlations that we need to be cognizant of: (a) default probability (PD) correlations across counterparties/issuers, (b) correlations of default, conditional on PDs, (c) correlation between PDs and loss given default (LGDs), and (d) correlation between credit risk and exposures, which are determined by the level of market risk (note that in derivative contracts the exposure comes counterparty failure on contracts that are in-the-money, and the extent of moneyness depends on the amount of price risk, computed as a mark-to-market value). When all these correlations are accounted for separately and collectively, the amount of risk capital to be maintained may be higher than currently stipulated.
4. Basel II does not specify how recovery rates will be specified to determine LGDs. We propose that the many techniques developed in the extant literature be examined to identify a possible approach to standardizing recovery rates across banks.
5. There is additional tail risk that comes from two sources. First, we cited evidence that LGDs are correlated in the cross-section of firms, resulting in contagion in losses on default. Second, we referred to empirical evidence in the literature that the single factor Gaussian risk model is rejected in favor of models with fatter tails. We recommend a careful evaluation of tail risk with a view to prevent being caught in systemic crises.
6. We have shown that longer VaR horizons lead to distortionary effects, especially when accounting for regimes and pro-cyclicality. We recommend that liquidity appropriate horizons be specified for each asset class.
7. The discussion shows that capital floors and minimum leverage requirements are valuable transitionary mechanisms, which may be enhanced with measures such as distance to default from the Merton (1974) model. The latter is recommended as a first step in bringing models to the implementation of the third pillar (market discipline) of the Basel II accord. This may eventually lead to risk-based capital computations in a top-down manner in the spirit of the Merton (1977) model.

Basel II represents a much needed improvement on Basel I, and consistently utilizes the technological advancements that have permeated the financial industry since Basel I. This paper has looked at some of the correlation issues that arise, and suggests that many of these issues may be addressed using the tools banks already employ for internal risk management.

Acknowledgements I am grateful for helpful discussions and input from an anonymous referee, George Chacko, Robert Jarrow, Paul Kupiec, Til Schuermann, and participants at the FDIC Conference on Banking.

References

Altman E, Brady B, Resti A, Sironi A (2005) The link between default and recovery rates: theory, empirical evidence and implications. J Bus 78:2203–2228
Artzner P, Delbaen F, Eber JM, Heath D (1999) Coherent measures of risk. Math Finance 9: 203–228, November
Research Task Force Concentration Risk Group of the Basel Committee on Banking Supervision (2004) Klaus Duellmann (chairman), BCBS Working Paper No. 15, November
Basel Committee on Banking Supervision (2005) International convergence of capital measurement and capital standards: a revised framework. http://www.bis.org/publ/bcbs118.htm, November
Buehler KS, D'Silva V, Pritsch G (2004) The business case for Basel II, vol 1 McKinsey Quarterly
Chorafas D (2004) Economic capital allocation with Basel II, Elsevier
Das S, Freed L, Geng G, Kapadia N (2006) Correlated default risk. J Fixed Income, Fall, 7–32
Das S, Geng G (2004) Correlated default processes: a criterion-based copula approach. J Investment Management 2(2):44–70
Das S, Duffie D, Kapadia N, Saita L (2007) Common failings: how corporate defaults are correlated. J Finance 62:93–117
Das S, Hanouna P (2006) Implied recovery. Working paper, Santa Clara University
Duffie D, Leandro S, Wang K (2007) Multi-period corporate default prediction with stochastic covariates. J Financ Econ v83(3):635–665
Engle R (2002) Dynamic conditional correlation—a simple class of multivariate GARCH Models. J Bus Econ Stat 20(3):339–350
Flannery M (2005) No pain, no gains? Effecting market discipline via 'Reverse Convertible Debentures'. In: Scott H (ed) Capital adequacy beyond basel, Oxford University Press, New York
Gordy M (2003) A risk-factor model foundation for ratings-based capital rules. J Financ Intermed 12:199–232
Gordy M (2004) Granularity adjustment in portfolio credit risk measurement. In: Giorgio Szëgo (ed) Risk measure for the 21st century. Wiley, New York
Gordy M, Lütkebohmert E (2006) Granularity adjustment for Basel II, Working Paper, Board of Governers of the Federal Reserve System
Gore G (2006) Correlation confusion. Risk, 59–61, July
Gup B (2004) The new basel capital accord. Thomson Publishing, New York
Jarrow R (2006) A critique of revised Basel II, Working Paper, Cornell University
Jeffery C, Chen X-L (2006) VaR breakdown. Risk, 43–48, July
Kupiec P (2005) Unbiased capital allocation in an asymptotic single risk factor (ASRF) model of credit risk, FDIC Working Paper
Kuritzkes A, Schuermann T, Weiner S (2003) Risk measurement, risk management, and capital adequacy in financial conglomerates. Brookings-Wharton Papers on Financial Services, pp 141–193
Lo A (2000) Risk management for hedge funds: introduction and overview. Working Paper, MIT
Merton RC (1974) On the pricing of corporate debt: the risk structure of interest rates. J Finance 29:449–470
Merton RC (1977) An analytic derivation of the cost of deposit insurance and loan guarantees: an application of modern option pricing theory. J Bank Finance 1:3–11
Merton RC, Perold AF (1993) Theory of risk capital in financial firms. J Appl Corp Finance 6(3): Fall, 16–32

Pan J, Singleton K (2005) Default and recovery implicit in the term structure of sovereign CDS spreads. Working Paper, Stanford University

Schuermann T (2004) What do we know about loss given default? In: Shimko D (ed) Credit risk models and management, 2nd edn. Risk Books Ch. 9, London, UK (February)

Unal H, Madan D, Guntay L (2001) Pricing the risk of recovery in default with APR violations. J Bank Finance 27:1001–1025

Zhang F (2003) What did the credit market expect of Argentina default? Evidence from default swap data. Working Paper, Federal Reserve Board