Nikola Tarashev
+41 61 280 9213
nikola.tarashev@bis.org

Haibin Zhu
+41 61 280 9164
haibin.zhu@bis.org

Measuring portfolio credit risk: modelling versus calibration errors[1]

A model-based assessment of credit risk is subject to both specification and calibration errors. Focusing on a well known credit risk model, we propose a methodology for quantifying the relative importance of alternative sources of such errors and apply this methodology to a large data set. We find that flawed calibration of the model can substantially affect the measured level of portfolio credit risk. By contrast, a model misspecification generally has a limited impact, especially for large, well diversified portfolios.

JEL classification: C15, G13, G21, G28.

In the wake of recent advances in risk management, models of portfolio credit risk have attracted increasing attention. The validation of such models is of interest to both market practitioners and supervisors, not least because errors in the measurement of credit risk could translate into errors in financial institutions' desired capital buffers. Such errors have different sources. They can be due to a violation of key modelling assumptions (ie *misspecification*) or to wrong estimates of key parameters (ie *flawed calibration*). Thus, a quantification of the relative importance of alternative sources of error in model outcomes would address issues of particular interest to the financial industry.

In this article, we tackle such issues in the context of the well known asymptotic single-risk factor (ASRF) model. This model implies that capital buffers can be set at the level of individual credit exposures and, thus, are *portfolio invariant*. This implication, which is incorporated in the internal ratings-based approach of the Basel II Framework, limits the data and operational requirements on users of the model. However, portfolio invariance rests critically on two assumptions, namely that the systematic component of credit risk is governed by only one common factor and that the portfolio is so finely grained that all idiosyncratic credit risk is diversified away. These assumptions have been criticised in the literature as being too strong and hence as being potential sources of misspecification errors.

[1] The views expressed in this article are those of the authors and do not necessarily reflect those of the BIS. The authors thank Marcus Jellinghaus for valuable help with the data.

By contrast, flawed calibration of the ASRF model, another potential source of error in the measurement of credit risk, has received comparatively little attention. This unbalanced focus in the literature overlooks the fact that accurate estimation of key parameters of this model imposes substantial data requirements. In fact, users who would be attracted by the model's simplicity, embodied in the "portfolio invariance" implication, are likely to also face difficulties in meeting those requirements.

We propose a general methodology for identifying different sources of misspecification and calibration errors in the measurement of portfolio credit risk. Using a data set that contains estimated probabilities of default (PDs) and asset return correlations for a large cross section of firms, we quantify the relative impact of such errors on outcomes of the ASRF model. Our illustrative exercise suggests that model-implied measures of portfolio credit risk are more sensitive to plausible calibration errors than to misspecification errors. This is especially true for larger portfolios where the violation of the granularity assumption is less pronounced.

The rest of this article is organised in four sections. After a brief overview of the related literature, the first section discusses at a conceptual level alternative sources of error in the calculation of portfolio credit risk. Then, the second section spells out the empirical methodology, which quantifies the relative importance of these sources in a unified framework. The third section describes the data and reports the empirical findings. The final section summarises the contribution of this analysis and identifies directions for future research.

Calculated versus target capital: conceptual issues

The ASRF model postulates that an obligor defaults when the value of its assets falls below a particular threshold. In addition, the model assumes that credit risk is related across obligors owing to the dependence (or the "loading") of their assets on a single common risk factor and that the portfolio consists of a large number of small exposures (ie is of fine "granularity"). In this model the capital that covers all portfolio losses with a desired probability can be calculated at the level of individual exposures. In turn, an exposure-specific capital depends solely on the corresponding PD and common-factor loading.

Model misspecification and flawed calibration …

The literature has paid closer attention to misspecification of the ASRF model than to the potentially flawed calibration of its inputs. Analysis of violations of ASRF assumptions has led to proposals of ways to mitigate the impact of such violations on capital calculations. The various proposals, which have been reviewed in BCBS (2006), attempt to strike a balance between the reduction of errors and the associated data or computational burden.[2] However, users who rely on the stylised ASRF model in order to

---

[2] Adjustments to capital measures that correct for violations of the ASRF granularity assumption have been derived in Martin and Wilde (2002), Emmer and Tasche (2003) and Gordy and Lütkebohmert (2006). In turn, violations of the single-common-factor assumption have been the focus of Pykhtin (2004), Düllmann (2006), Garcia Cespedes et al (2006) and Düllmann and Masschelein (2006). In addition, Heitfield et al (2006) and Düllmann et al (2006)

alleviate such a burden are also likely to face constraints in estimating the model's parameters. Thus, they are prone to an imperfect calibration of the model, which would generate additional bias in the assessment of portfolio credit risk.[3]

… affect popular measures of portfolio credit risk

In this special feature, we extend previous analyses by examining four sources of error in a model-based assessment of portfolio credit risk. Two of these sources relate to the possible misspecifications of the ASRF model that have received much attention in the literature. The other two relate to an erroneous calibration of the correlation of credit risk across exposures.[4]

In the remainder of this section, we define and discuss each of these sources of error at a conceptual level. The metric we use to compare different model outcomes is credit value-at-risk (net of expected losses), which is equivalent to the capital buffer necessary to cover default losses with a desired probability.[5] The comparison is based on a hypothetical benchmark assessment of risk, which assumes knowledge of all relevant parameters and is referred to as "target capital", and alternative assessments, which are subject to one or more of the above-mentioned sources of error.

Multiple factors of credit risk

The obligors in a credit portfolio are likely to be affected not only by aggregate economic conditions but also by conditions relevant for specific business lines. Conceptually, these various conditions can be summarised in mutually independent and often unobservable risk factors. If several of these factors are of material importance, the single-factor assumption of the ASRF model would be violated. This would entail systematic errors in model-implied measures of credit risk and, consequently, in the capital set aside to compensate for it.

Violations of the single-factor assumption …

A violation of the single-factor assumption is conceptually different from a failure to measure the impact of multiple factors on the correlation across obligors. Such a failure is independent of a modelling misspecification and can arise, for example, when higher concentration in a particular industrial sector is not captured in the estimated average correlation. However, even if the average correlation across obligors is measured accurately, an erroneous single-factor assumption ignores the fact that there are multiple sources of default clustering. This leads to an underestimation of the probability of a large number of defaults and, consequently, to an underestimation of the target capital.

---

have examined both assumptions in the context of representative portfolios of US and European banks, respectively.

3 For a discussion of the impact of estimation errors on capital calculations, see Löffler (2003).

4 In order to focus on "pure portfolio" characteristics of credit losses, we abstract from errors in the measurement of PDs and losses-given-default (LGDs).

5 In this article the terms "model outcome", "assessment of portfolio risk" and "capital" are used interchangeably. Importantly, our capital calculations do not correspond to "regulatory capital", which reflects considerations of bank supervisors, or to "economic capital", which reflects additional strategic and business objectives of financial firms.

Graph 1 (left-hand panel) illustrates such an effect in the context of a stylised portfolio, in which exposures are the same across obligors and have homogeneous PDs and losses-given-default (LGDs). In this portfolio, a fraction ω of the obligors belong to group 1, while 1–ω belong to group 2. The associated asset returns are uncorrelated across groups but are affected by a group-specific common factor. Thus, increasing the value of ω between 0 and 1 increases the relative importance of the first common factor for the overall credit risk in the portfolio. The red line plots the target capital as a function of ω,[6] while the green line portrays an alternative capital calculation implied by a single-factor structure. This structure matches exactly the average level of asset return correlations across exposure pairs.

The difference between the red and green lines illustrates the multi-factor effect that we analyse empirically below. This difference is largest when the role of multiple factors is greatest and, hence, when a single-factor structure approximates most poorly the dispersion of asset return correlations. In our example, this occurs at ω = 1/2.

Modelling misspecifications and capital, illustrative examples

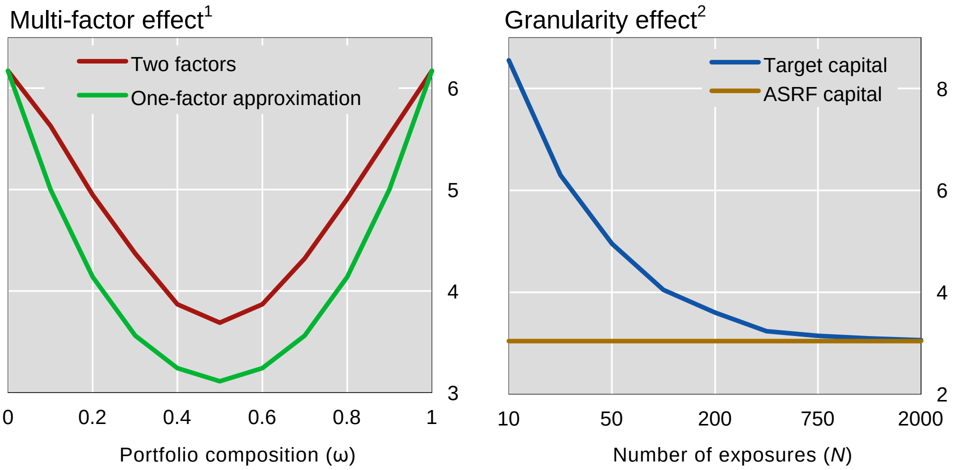

Note: Calculated capital, in per cent, per unit of aggregate exposure, on the vertical axes. For each panel, PD = 1% and LGD = 45% are the same across exposures.

[1] The portfolio consists of two groups of exposures, with ω denoting the weight of the first group. Within each group, the asset return correlation equals 20% for all exposure pairs. The inter-group correlation is zero. The red line plots the target capital level, which incorporates two common factors. The green line plots the capital calculated under a one-common-factor approximation of the correlation structure. This approximation imposes the same common-factor loadings across firms and does not change the average asset return correlation. [2] The blue line plots target capital for a portfolio in which all pairwise asset return correlations equal 10%. The number of exposures in this portfolio varies across the horizontal axis. The brown line plots the corresponding capital estimates when *N* is assumed to be infinite.

Graph 1

[6] Target capital is lowest at ω = 1/2 because at this value the portfolio is evenly diversified between the two common factors, which minimises the probability of large losses. The dependence of the desired capital buffer on the relative exposure to multiple factors is studied by Düllmann and Masschelein (2006), at both a theoretical and an empirical level.

Granularity

… and the granularity assumption …

The “perfect granularity” assumption of the ASRF model postulates that all exposure-specific, ie idiosyncratic, risk is diversified away. Since this cannot be attained by any real-world portfolio, the granularity assumption leads to an underestimation of the overall credit risk and, consequently, to an underestimation of target capital. Given an overall correlation across exposures, this underestimation is smaller for a portfolio that comprises more obligors and, thus, benefits from greater diversification gains. This is illustrated in the right-hand panel of Graph 1.

Correlation level

Even if the ASRF model is correctly specified, calculated capital may be affected by errors in the inputs to this model. For instance, a user of the model who has data constraints may choose to rely on readily available external estimates of the asset return correlations in popular credit indices. Such estimates would lead to a discrepancy between target and calculated capital to the extent that the underlying indices were not representative of the user’s own portfolio.

… and erroneous correlation estimates …

One driver of the potential discrepancy is a bias in the average level of the asset correlations underpinning the calculation of capital. When this bias is positive, for instance, it inflates the probability that a large number of defaults might occur simultaneously and leads to an overestimation of the target capital. Conversely, a negative bias leads to an underestimation. This is the correlation level effect we examine empirically below.

Dispersion in pairwise correlations

The effect of a bias in the overall level of calibrated correlations could be augmented by errors in the dispersion of these correlations across exposure pairs. Such errors are likely to emerge either if users of the ASRF model rely on external estimates of asset return correlation or if they apply the average of internal estimates to all exposures. This would lead to a correlation dispersion effect on calculated capital. Even though in practice this effect is likely to be tightly related to the correlation level effect, our empirical methodology will disentangle the two in order to quantify their separate roles.

The effect of correlation dispersion on calculated capital is seen most clearly in a stylised example. Suppose that all firms in one portfolio have homogeneous PDs and exhibit homogeneous pairwise asset return correlations. Suppose further that a second portfolio is characterised by the same PDs and average asset return correlation but includes a group of firms that are more likely to default together. The second portfolio is more likely to experience several simultaneous defaults and, thus, requires higher capital in order to attain solvency with the same probability.

However, this result can be weakened or even reversed if PDs vary across firms. To see why, suppose that the strongly correlated firms in the second portfolio are the ones that have the lowest individual PDs. In other words, the firms that are likely to generate multiple defaults are less likely to default. This may lower the probability of default clustering, depressing the target capital level below that for the first portfolio.

Empirical methodology

Our empirical methodology comprises two general steps. In the first step, we construct a large (small) hypothetical portfolio that comprises equal exposures to 1,000 (200) firms.[7] The sectoral composition of this portfolio is designed to be in line with the typical loan portfolio of large US wholesale banks.[8] Given the constraints of such a composition, the portfolio is drawn at random from our sample of firms. Since each draw could be affected by sampling errors, we examine 3,000 different draws for both large and small portfolios.

… can be analysed in a unified framework …

For a portfolio constructed in the first step, the second step calculates five alternative capital measures, which differ in the underlying assumptions regarding the interdependence of credit risk across exposures.[9] Each of these alternatives employs the same set of PD estimates, and assumes that LGD equals 45% for all exposures and that asset returns are normally distributed.[10] We order the measures so that each measure differs from a previous one owing to a *single* assumption.[11]

The first measure is the *target* capital, which incorporates data on asset return correlation estimates. Assuming that these and the other risk parameter estimates we adopt (as well as our distributional assumption) are accurate, we conduct a Monte Carlo simulation to construct the probability distribution of default losses at the one-year horizon. Then, we set the target capital to a level that covers unexpected default losses with a probability of 99.9%,[12] recognising that our methodology also applies to alternative definitions of target capital.

… as sources of deviations from target capital

The second capital measure differs from the target one owing to a restriction on the number of common factors governing asset returns. Specifically, we use a correlation matrix of asset returns that can emerge in the presence of a single common risk factor but fits as closely as possible the

---

7 In this analysis, the distinction between large and small portfolios is not based on the size of the aggregate exposure but reflects different degrees of diversification across individual exposures.

8 Such a portfolio does not incorporate consumer loans and, thus, may not fully represent all aspects of credit risk. To construct a large portfolio, we apply the 40 sector-specific weights provided by Heitfield et al (2006). For a small portfolio, we rescale the 10 largest sectoral weights so that they sum up to one and set all remaining weights to zero.

9 The box on page 89 provides further detail on the calculation of each of the five capital measures.

10 The particular data used in this article are described in the next section.

11 We choose one of several possible orderings of the five capital measures. On the basis of background analysis, we are confident that an alternative ordering would not change our main conclusions significantly, even if it altered specific numerical results.

12 The covered level of unexpected losses equals the 99.9th percentile of the distribution of credit losses minus the mean of this distribution.

Calculating capital measures: technical details

This box outlines three general methods for calculating the distribution of default losses. These methods are used to derive the five capital measures considered in the text.

The first method relies on Monte Carlo simulations and delivers the target capital level. This method can be applied to any portfolio comprising N equally weighted exposures, provided that the associated probabilities of default, PD, losses-given-default, LGD, and correlation matrix of asset returns, R, are known. The method consists of three general steps. In the first step, one uses the vector of PDs and the assumption that asset returns are distributed as standard normal variables to obtain an N x 1 vector of default thresholds. In the second step, one draws an N x 1 vector from N standard normal variables whose correlation matrix is R. The number of entries in this vector that are smaller than the corresponding default threshold is the number of simulated defaults for the particular draw. In the third step, one repeats the second step a large number of times to derive the probability distribution of the number of defaults. Denoting this distribution's (1 – α)th percentile by β and the average PD in the portfolio by $A(PD)$, the target capital for a credit value-at-risk confidence level of (1 – α) equals $LGD^*(\beta - A(PD))$ per unit of exposure.[①]

The second method relies on the so-called Gaussian copula, which is outlined in detail in Gibson (2004). This method delivers the second capital measure in the main text and rests on the assumptions that (i) the portfolio consists of N equally weighted exposures with identical LGD[②] and (ii) only one common factor drives credit risk. To apply this method, one needs to calibrate the LGD, obtain values for the firm-specific PDs and estimate firm-specific loading coefficients, l_i , which are defined by the following equation:

$$V_i = l_i M + \sqrt{1 - l_i^2} Z_i \qquad (1)$$

where V_i is the asset value of firm i, M is the common risk factor and Z_i is the idiosyncratic risk factor. Equation (1) implies that the correlation between the asset returns of firm i and j equals $l_i{*}l_j$. A particular estimate of l_i is obtained by fitting the single-common-factor assumption to the original correlation structure in a mean-squared-error sense, ie by minimising the following sum:

$$\sum_{i=1}^{N-1} \sum_{j>i} \left(\rho_{ij} - l_i l_j\right)^2 \qquad (2)$$

where ρ_{ij} is an element of the correlation matrix R. Andersen et al (2003) provide an efficient algorithm for solving this minimisation problem. Estimated in this way, the loading coefficients l_i account almost exactly for the average pairwise correlation in R.

The third method, which applies to the other three measures in the main text, is a special case of the ASRF model. In comparison to the second method, it makes the additional assumption that all idiosyncratic risk is diversified away. This implies that the capital buffer for the portfolio is the sum of exposure-specific capital buffers, κ_i, which are calculated as follows:

$$\kappa_i = LGD * \left[\Phi\left(\frac{\Phi^{-1}(PD_i) - l_i \Phi^{-1}(\alpha)}{\sqrt{1 - l_i^2}} \right) - PD_i \right] \qquad (3)$$

---

① This article sets N = 200 or 1,000, LGD = 45%, α = 0.1% and PD and R as estimated by Moody's KMV. In addition, the third step of the first method carries out 500,000 Monte Carlo simulations. ② For a calculation of the probability distribution of credit losses when individual exposures have different LGDs and portfolio weights, see Hull and White (2004).

unrestricted correlation matrix underpinning the target capital calculation. The difference between the resulting capital estimate and the target level is denoted the "*multi-factor effect*".

The third measure differs from the second one in that it assumes, in addition, that all idiosyncratic credit risk is diversified away. In other words, it ignores the impact of imperfect granularity in the portfolio. This assumption allows one to apply the ASRF model, which delivers an analytic solution for capital calculations. The difference between the third and second measures is the "*granularity effect*".

The fourth measure differs from the third one in that it is based on the assumption that asset return correlations are the same across all exposures in the portfolio. The common correlation, which is set equal to the average of the pairwise correlations underpinning the third capital measure, is used as an input to the ASRF model. The difference between the fourth and third capital measures is denoted the "*correlation dispersion effect*".

Finally, the fifth measure incorporates a bias in the estimates of asset return correlations. This measure differs from the fourth one in that it relies on a standard rule-of-thumb value of the (common) asset return correlation. The difference between the fifth and fourth measures is denoted the "*correlation level effect*".

An important feature of this methodology is that it allows one to quantify the relative importance of alternative drivers of capital miscalculations. Specifically, the methodology can be applied to dissect the difference between the fifth capital measure, which we henceforth dub the "shortcut" one, and the first, ie target, capital measure. By construction, this difference equals exactly the sum of the multi-factor, granularity, correlation dispersion and correlation level effects.

Empirical results

In this section, we employ the methodology outlined above to investigate the discrepancy between target and shortcut capital for simulated portfolios with realistic features. In addition, we derive the relative importance of alternative drivers of this discrepancy.

Data

Data on risk parameter estimates suggest that …

Our empirical analysis relies on two data sets provided by the commercial service Moody's KMV. One data set consists of one-year expected default frequencies (EDF™) that are point-in-time estimates of individual PDs, while the other comprises estimates of pairwise asset return correlations implied by the GCorr model. Both EDF and GCorr models are based on an extended and operational version of the seminal framework of Merton (1974), which is broadly consistent with the ASRF model.[13]

Combining these two data sets, we obtain a pool of 10,891 non-financial firms, for which individual PD and pairwise correlation estimates are available.

[13] See Crosbie and Bohn (2003), Das and Ishii (2001) and Crosbie (2005) for a description of these proprietary models and related references. The Moody's KMV sample comprises only firms with publicly traded equities. A multi-factor loading structure is employed for the estimation of the GCorr model.

Characteristics of simulated portfolios[1]

In per cent

A. Large portfolios (1,000 exposures)	Mean	Standard deviation	Median	Minimum	Maximum
Average PD	2.42	0.19	2.42	1.79	3.12
Standard deviation of PDs	5.16	0.26	5.16	4.25	6.14
Median PD	0.26	0.03	0.26	0.18	0.36
Average correlation[2]	9.78	0.22	9.77	9.14	10.73
Standard deviation of loadings[3]	9.33	0.31	9.32	8.33	10.47
Corr (PD, loadings)[4]	–20.00	2.04	–20.10	–26.70	–12.80
B. Small portfolios (200 exposures)	**Mean**	**Standard deviation**	**Median**	**Minimum**	**Maximum**
Average PD	2.28	0.36	2.26	1.24	3.68
Standard deviation of PDs	5.05	0.53	5.06	3.01	6.89
Median PD	0.24	0.05	0.23	0.11	0.55
Average correlation[2]	10.49	0.44	10.48	8.99	12.00
Standard deviation of loadings[3]	10.54	0.70	10.55	7.80	12.79
Corr (PD, loadings)[4]	–19.80	4.59	–20.20	–31.80	–1.20

Note: The calculations in this table use 3,000 simulated portfolios for each portfolio size and are carried out in two steps. First, portfolio-specific statistics specified by row headings are calculated for each simulated portfolio. Second, summary statistics specified by column headings are calculated for each of the portfolio-specific statistics calculated in the first step.

[1] Based on Moody's KMV estimates of PDs and asset return correlations for July 2006. [2] Based on all pairwise correlations. [3] The derivation of common-factor loadings assumes that there is a single common factor and implements the procedure outlined in the box (page 89). [4] The sample correlation between PDs and loadings on the single common factor.

Table 1

The vast majority of the firms in the sample are headquartered either in the United States (52% of the total number) or in western Europe (40%). For illustrative purposes, we use EDF and GCorr estimates in July 2006 as the "true" PDs and correlations underpinning the target capital level. Of course, any error in these estimates would warrant a revision of the target capital.

Table 1 summarises the characteristics of the simulated portfolios. For both large and small portfolios, the distribution of EDFs has a long right tail, with the median values much lower than the mean. Correlation estimates are clustered mainly between 5 and 25%, with their mean standing at 9.78% for large portfolios and 10.49% for small ones. In addition, reflecting the benign credit conditions during the sample period, more than 10% of the sample firms have the lowest PD estimates permitted by the EDF model (ie 0.02%).

Target versus shortcut capital levels

… deviations from target capital can be large …

We proceed to quantify and decompose the difference between target and shortcut capital measures (Table 2). For illustrative purposes, the constant correlation underlying shortcut calculations is set at 12%. This is about

Four sources of error in estimated capital[1]

Per unit of aggregate exposure, in per cent

A. Large portfolios (1,000 exposures)					
	Mean	Standard deviation	Median	95% interval	50% interval
Target capital[2]	2.95	0.16	2.95	[2.64, 3.27]	[2.84, 3.05]
Deviation from the target due to:[3]					
Multi-factor effect	–0.03	0.03	–0.05	[–0.09, 0]	[–0.05, 0]
Granularity effect	–0.11	0.01	–0.11	[–0.14, –0.09]	[–0.12, –0.10]
Correlation dispersion effect	0.35	0.04	0.35	[0.27, 0.43]	[0.32, 0.38]
Correlation level effect	0.55	0.06	0.55	[0.44, 0.66]	[0.52, 0.59]
"Shortcut" capital (correlation = 12%)	3.71	0.18	3.71	[3.37, 4.06]	[3.59, 3.83]
Memo: correlation level effect if:					
Correlation = 6%	–0.96	0.07	–0.96	[–1.11, –0.83]	[–1.00, –0.91]
Correlation = 18%	2.01	0.09	2.01	[1.84, 2.18]	[1.95, 2.07]
Correlation = 24%	3.47	0.13	3.47	[3.23, 3.72]	[3.39, 3.56]
B. Small portfolios (200 exposures)					
	Mean	Standard deviation	Median	95% interval	50% interval
Target capital[2]	3.35	0.30	3.34	[2.78, 3.94]	[3.15, 3.53]
Deviation from the target due to:[3]					
Multi-factor effect	–0.04	0.10	0	[–0.23, 0]	[0, 0]
Granularity effect	–0.53	0.07	–0.53	[–0.65, –0.41]	[–0.59, –0.47]
Correlation dispersion effect	0.38	0.11	0.37	[0.17, 0.58]	[0.30, 0.45]
Correlation level effect	0.36	0.11	0.36	[0.15, 0.61]	[0.29, 0.44]
"Shortcut" capital (correlation = 12%)	3.52	0.34	3.51	[2.85, 4.23]	[3.28, 3.75]
Memo: correlation level effect if:					
Correlation = 6%	–1.07	0.12	–1.07	[–1.31, –0.85]	[–1.15, –0.99]
Correlation = 18%	1.76	0.19	1.75	[1.41, 2.15]	[1.63, 1.88]
Correlation = 24%	3.15	0.27	3.14	[2.65, 3.70]	[2.97, 3.33]

[1] Summary statistics for the simulated portfolios underpinning Table 1 (3,000 for each portfolio size). The column entitled "95% interval" reports the 2.5th and 97.5th percentiles of the statistics specified in the particular row heading. The column entitled "50% interval" reports the corresponding 25th and 75th percentiles. [2] Based on Moody's KMV estimates of PDs and asset return correlations and a Monte Carlo procedure for calculating the probability distribution of default losses. [3] Four sources of deviation from the target capital level; a negative sign implies underestimation. The sum of the target capital level and the four deviations equals the shortcut capital level. Each deviation is based on the assumptions underlying previous deviations plus one additional assumption: (a) for the multi-factor effect, the correlation matrix underpinning the target capital level is approximated under the assumption that there is a single common factor; (b) for the granularity effect, there is the additional assumption that the number of firms is infinite; (c) for the correlation dispersion effect, the additional assumption is that the loadings on the single common factor are the same across exposures; (d) for the correlation level effect, the additional assumption imposes a different level for the constant pairwise correlation. See the box on page 89 for further detail on alternative capital measures. Table 2

2 percentage points higher than the average asset return correlation in the simulated portfolios (recall Table 1) and is close to the 12.5% rule-of-thumb correlation suggested by Lopez (2004).

The results show that the shortcut capital measure can be significantly higher than the corresponding target level. The difference is much more pronounced in the context of large portfolios, for which it amounts on average to 76 basis points (per unit of exposure, or 26% of the target capital). This is predominantly the result of the correlation dispersion and correlation level effects. By contrast, these two effects are almost fully offset by the granularity effect in small portfolios, for which the shortcut capital estimate is 5% higher than the target level.[14] The following subsections discuss in some detail the four alternative effects behind the overall discrepancies between target and shortcut capital.[15]

Multi-factor effect

The multi-factor effect lowers the model-implied capital measure, for the reasons outlined above, but its quantitative impact is almost negligible. Imposing a single-factor structure on asset returns leads to a capital allocation on large portfolios that is, on average, 1% lower than the target level. At 1.2%, this decline is only slightly larger for small portfolios.

The low importance of the multi-factor effect is a result of the fact that a single-factor framework approximates quite well the multi-factor structure in the data. For the portfolios used in this exercise, the single-factor framework outlined in the box on page 89 explains almost perfectly the level of the original correlations (with an error of less than 2 basis points) and accounts for the bulk (76% on average) of the cross-sectional variation in pairwise correlations. The robustness of this finding to alternative portfolio specifications and alternative estimates of risk parameters is an important question for future research.

Granularity effect

As expected, the granularity assumption of the ASRF model leads to an underestimate of the target capital ratio. Importantly, the underestimation increases when the size of the portfolio decreases. As Table 2 reports, the granularity effect leads to a 4% underestimation of the target capital for large, diversified portfolios and a 16% underestimation for small, less diversified portfolios. These results are in line with previous analyses of the granularity effect.[16]

In practice, the size of the exposures would vary across obligors, which would complicate the analysis of the granularity effect. For example, a portfolio that consists of a large number of exposures but is highly concentrated in a subset of them can be associated with a larger granularity effect than a portfolio with a smaller number of equally weighted exposures. Our

---

[14] Even though the exercise focuses on a particular sectoral distribution of exposures, credit risk does differ across the simulated portfolios. Accordingly, columns 2 to 5 in Table 2 report descriptive statistics of the distribution of the portfolio-specific capital estimates.

[15] In quantifying the magnitude of each effect, the adopted sign is such that an effect can be added to the target capital level or subtracted from the corresponding shortcut level.

[16] For instance, Gordy and Lütkebohmert (2006) propose an adjustment formula to correct for the granularity effect. An application of this formula (equation (6) in their paper) matches exactly a granularity effect that leads to a 5.4% underestimate of the target capital for large portfolios and a 24% underestimate for small portfolios.

methodology could also accommodate such cases, but we abstract from them in this special feature in order to simplify the exposition.[17]

Correlation dispersion effect

… especially if the model is erroneously calibrated

Equalising the asset return correlations across exposure pairs causes calculated capital to be more conservative than the target level. In particular, target capital is overestimated by 12% for large portfolios. At 11% for small portfolios, this overestimation classifies the correlation dispersion effect as the most important of the four considered sources of discrepancies between shortcut and target capital. In line with the intuition presented above, the positive sign of the correlation dispersion effect is due to the fact that, in our sample, higher-PD exposures tend to be less correlated among themselves (Table 1).[18]

Correlation level effect

A mismatch between the average correlations underpinning target and shortcut capital calculations would also have a substantial effect. Increasing the average asset return correlation from 9.8% (the level estimated by Moody's KMV for the simulated portfolios) to 12% leads to a 19% overestimation of the target capital for large portfolios. For small portfolios, an average asset return correlation of 12% implies a smaller but still significant overestimation of 11%.

This result is not surprising, because a higher average asset return correlation translates into a higher probability of default clustering, which raises the estimated capital. Alternatively, however, the average level of asset return correlations may be underestimated, which would lead to insufficient capital. Table 2 reports that setting this level to 6% would lead to underestimating the target capital level by about 32% for both portfolios.[19]

Conclusion

In this article, we developed an approach to evaluating errors in the measurement of portfolio credit risk. In particular, we used this approach to quantify the magnitude of different sources of a discrepancy between a predefined target capital level and a shortcut alternative, which is based on the ASRF model and rule-of-thumb correlation estimates. On the basis of simulated

---

[17] Accommodating disparate exposures would introduce an additional dimension in portfolio characteristics, requiring the simulation of a greater variety of hypothetical portfolios and making it more difficult to interpret the multi-factor and correlation dispersion effects.

[18] The negative relationship between PDs and correlations (ie loading coefficients) is likely to be a general phenomenon. For example, Dev (2006) finds that global factors often play bigger roles for firms of better credit quality.

[19] Table 2 reports the correlation level effect based on alternative levels of average correlations: 6%, 18% and 24%. These alternatives correspond to rule-of-thumb correlation values reported in previous studies (between 5 and 25%) and to plausible estimation errors. As regards such errors, Tarashev and Zhu (2007) show that, for a true constant correlation of 9.78% and five years of monthly data on asset returns, the 95% confidence interval for the estimated average correlation is between 6.4 and 13.3%.

portfolios, we found that plausible errors in estimated asset return correlations could lead to substantial deviations from the target capital levels for both large and small portfolios. By contrast, the violation of key assumptions of the ASRF model, ie the single-factor or the perfect granularity assumption, tend to result in relatively smaller errors in calculated capital. The only exception is that the granularity assumption does have a significant impact for small portfolios.

The illustrative nature of our analysis identifies different avenues for future research. For one, it would be valuable to analyse the robustness of our empirical results to alternative portfolio specifications and to different (realistic) values of PDs, LGDs and asset return correlations. In addition, it would be important to derive rigorously the range of plausible estimation errors in the parameters used to calculate portfolio credit risk and to study the implications of alternative assumptions as regards the distribution of asset returns. Tarashev and Zhu (2007) attempt to address this latter set of issues.

References

Andersen, L, J Sidenius and S Basu (2003): "All your hedges in one basket", *Risk*, pp 67–72.

Basel Committee on Banking Supervision (2006): "Studies on credit risk concentration: an overview of the issues and a synopsis of the results from the Research Task Force project", *Basel Committee on Banking Supervision Working Paper*, no 15.

Crosbie, P (2005): "Global correlation factor structure: modelling methodology", *Moody's KMV Documents*.

Crosbie, P and J Bohn (2003): "Modelling default risk", *KMV White Paper*.

Das, A and S Ishii (2001): "Methods for calculating asset correlations: a technical note", *KMV Documents*.

Dev, A (2006): "The correlation debate", *Risk*, October.

Düllmann, K (2006): "Measuring business sector concentration by an infection model", *Deutsche Bundesbank Discussion Paper*.

Düllmann, K and N Masschelein (2006): "Sector concentration in loan portfolios and economic capital", *Deutsche Bundesbank Discussion Paper*.

Düllmann, K, M Scheicher and C Schmieder (2006): "Asset correlations and credit portfolio risk – an empirical analysis", working paper.

Emmer, S and D Tasche (2003): "Calculating credit risk capital charges with the one-factor model", *Journal of Risk,* vol 7, pp 85–101.

Garcia Cespedes, J C, J A de Juan Herrero, A Kreinin and D Rosen (2006): "A simple multi-factor "factor adjustment" for the treatment of credit capital diversification", *Journal of Credit Risk*, vol 2, pp 57–85.

Gibson, M (2004): "Understanding the risk of synthetic CDOs", *Federal Reserve Board Finance and Economics Discussion Series*.

Gordy, M and E Lütkebohmert (2006): "Granularity adjustment for Basel II", *Deutsche Bundesbank Working Paper*, forthcoming.

Heitfield, E, S Burton and S Chomsisengphet (2006): "Systematic and idiosyncratic risk in syndicated loan portfolios", *Journal of Credit Risk*, vol 2, pp 3–31.

Hull, J and A White (2004): "Valuation of a CDO and an n-th to default CDS without Monte Carlo simulation", *Journal of Derivatives*, vol 12, pp 8–23.

Löffler, G (2003): "The effects of estimation error on measures of portfolio credit risk", *Journal of Banking and Finance*, vol 27, pp 1427–53.

Lopez, J (2004): "The empirical relationship between average asset correlation, firm probability of default and asset size", *Journal of Financial Intermediation*, vol 13, pp 265–83.

Martin, R and T Wilde (2002): "Unsystematic credit risk", *Risk*, pp 123–28.

Merton, R (1974): "On the pricing of corporate debt: the risk structure of interest rates", *Journal of Finance*, vol 29, pp 449–70.

Pykhtin, M (2004): "Multi-factor adjustment", *Risk*, pp 85–90.

Tarashev, N and H Zhu (2007): "Modelling and estimation errors in measures of portfolio credit risk", *BIS Working Papers*, forthcoming.