Specification and Calibration Errors in Measures of Portfolio Credit Risk: The Case of the ASRF Model*

Nikola Tarashev and Haibin Zhu
Bank for International Settlements

This paper focuses on the asymptotic single-risk-factor (ASRF) model in order to analyze the impact of specification and calibration errors on popular measures of portfolio credit risk. Violations of key assumptions of this model are found to be virtually inconsequential, especially for large, well-diversified portfolios. By contrast, flaws in the calibrated interdependence of credit risk across exposures, caused by plausible small-sample estimation errors or rule-of-thumb values of asset return correlations, can lead to significant inaccuracies in measures of portfolio credit risk. Similar inaccuracies arise under standard assumptions regarding the tails of the distribution of asset returns.

JEL Codes: G21, G28, G13, C15.

*We are indebted to Michael Gordy and Kostas Tsatsaronis for their encouragement and useful comments on an earlier version of this paper. We thank Klaus Düllmann, Wenying Jiangli, and participants at a BIS workshop, the 3rd International Conference on Credit and Operational Risk at HEC Montréal, and the 17th Annual Derivatives Securities and Risk Management Conference at the FDIC for insightful suggestions. We are also grateful to Marcus Jellinghaus for his valuable help with the data. The views expressed in this paper are our own and do not necessarily represent those of the Bank for International Settlements.
Author contact: Nikola Tarashev: Research and Policy Analysis, Bank for International Settlements, Basel, Switzerland; Tel.: 41-61-280-9213; Fax: 41-61-280-9100; E-mail: nikola.tarashev@bis.org. Haibin Zhu: Research and Policy Analysis, Bank for International Settlements, Basel, Switzerland; Tel.: 41-61-280-9164; Fax: 41-61-280-9100; E-mail: haibin.zhu@bis.org.

1. Introduction

Assessments of portfolio credit risk have attracted much attention in recent years. One reason is that participants in the increasingly popular market for structured finance products rely heavily on estimates of the interdependence of credit risk across various exposures.[1] Such estimates are also of principal interest to financial supervisors who, in enforcing new standards in the banking and insurance industries, have to ensure that regulatory capital is closely aligned with credit risk.

This paper investigates the well-known asymptotic single-risk-factor (ASRF) model of portfolio credit risk, which underpins the internal-ratings-based (IRB) approach of the Basel II framework (Basel Committee on Banking Supervision 2005). The popularity of this model stems from its implication that the contribution of each exposure to the credit value-at-risk (VaR) of the portfolio—defined as the maximum default loss that can be incurred with a given probability over a given horizon—is independent of the characteristics of the other exposures. This implication has been derived rigorously in Gordy (2003) and is known as *portfolio invariance* of marginal credit VaR contributions. The implication has been interpreted as alleviating the data requirements and computational burden on users of the model. Indeed, portfolio invariance implies that credit VaR can be calculated solely on the basis of exposure-specific parameters, including individual probability of default (PD), loss given default (LGD), and dependence on the common factor.

Its popularity notwithstanding, the "portfolio invariance" implication of the ASRF model hinges on two strong assumptions that have been criticized as sources of *specification errors*. Namely, the model assumes that the systematic component of credit risk is governed by a single common factor and that the portfolio is so finely grained that all idiosyncratic risks are diversified away. Violations of the "single-factor" and "perfect-granularity" assumptions would translate directly into erroneous assessments of portfolio credit risk.

[1] Examples of structured finance products are collateralized debt obligations (CDOs), nth-to-default credit default swaps (CDSs), and CDS indices.

Moreover, putting specification issues aside, a practical implementation of the ASRF model may be quite challenging. An important reason is that the portfolio invariance implication does *not* allow a user of the model to consider any particular exposure in isolation. Namely, estimates of exposure-specific dependence on the common factor, which are required for a calibration of the ASRF model, hinge on information about the correlation structure in the portfolio or about the common factor itself.[2] When such portfolio- or market-wide information is imperfect, the user will implement a *flawed calibration* of the model, which will be another source of errors in measured portfolio credit risk.

A contribution of this paper is to develop a *unified* method for quantifying the importance of model specification and calibration errors in assessments of portfolio credit risk. In order to implement this method, we rely on a large data set that comprises Moody's KMV estimates of PDs and pairwise asset return correlations for nearly 11,000 nonfinancial corporations worldwide. We use these estimates as the actual credit-risk parameters of hypothetical portfolios that are designed to match the industrial-sector concentration of typical portfolios of U.S. wholesale banks. For each hypothetical portfolio, we derive the "true" probability distribution of default losses and then condense this distribution into unexpected losses, which are defined as a credit VaR net of expected losses. This summary statistic is equivalent to a "target" capital measure necessary to cover default losses with a desired probability.[3] Target capital can be compared directly to a "shortcut" capital measure, which relies on the ASRF model and incorporates a rule-of-thumb calibration of the interdependence of credit risk across exposures.

[2]The IRB capital formula of Basel II abstracts from this calibration issue by postulating that firm-specific dependence on the single common factor is determined fully by the level of the corresponding PD.

[3]In this paper, we use the terms "assessment of portfolio credit risk" and "capital measure" interchangeably. Importantly, our capital measures 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.

We decompose the difference between the target and shortcut capital measures into four non-overlapping and exhaustive components. Two of these components, which we attribute to a "multi-factor" effect and a "granularity" effect, relate to misspecification of the ASRF model. The other two components, which we derive after transforming the correlation structure to be consistent with the ASRF model, relate to errors in the calibration of the interdependence of credit risk across exposures.[4] Specifically, the calibration errors we consider arise either from an overall bias in the measured correlations of firms' asset returns—which gives rise to what we dub a "correlation level" effect—or from noise in the measured dispersion of these correlations across pairs of firms—a "correlation dispersion" effect.

Another contribution of this paper is that it provides two additional perspectives on flaws in the calibration of the ASRF model. First, we calculate deviations from a desired capital buffer that arise *not* from the adoption of rule-of-thumb parameter values but from plausible small-sample errors in asset return correlation *estimates.* Second, motivated by the analysis in Gordy (2000) and Frey and McNeil (2003), we examine the importance of errors in the calibration of tail dependence among asset returns. The impact of such errors on capital measures is similar to but materializes independently of the impact of errors in estimated asset return correlations.

Our conclusion is that errors in the practical implementation, as opposed to the specification, of the ASRF model are the main sources of potential miscalculations of credit risk in large portfolios. Specifically, the misspecification-driven multifactor effect leads a user of the model to underpredict target capital buffers by only 1 percent. This is because a single-factor approximation, if chosen optimally, fits well the correlation structure of asset returns in our data. In addition, the granularity effect results in a 5 percent underprediction of the target level. In comparison, assessments of portfolio credit risk are considerably more sensitive to possible miscalibrations of the single-factor model. The correlation dispersion effect, for example, leads to capital measures that are 12 percent higher than the target measure. In turn, the correlation level effect causes a roughly 8 percent

[4] In order to sharpen the analysis, we do not analyze the implications of errors in the estimates of PDs and LGDs.

overprediction (underprediction) of the target measure for each percentage point of positive (negative) error in the average correlation coefficient. Furthermore, plausible small-sample errors in correlation estimates—arising when users of the model have five to ten years of monthly asset returns data—translate into capital measures that may deviate from the target level by 30 to 45 percent. Finally, data on asset returns are at odds with the conventional multinormality assumption. Specifically, this assumption implies tail-of-distribution dependence among asset returns that is too low and translates into an underestimation of the target capital by 22 to 86 percent.

Among the four effects on capital measures, only the granularity effect is sensitive to the number of exposures in the portfolio. When this number decreases, the portfolio maintains a larger portion of idiosyncratic risks. Thus, we are not surprised to find that the granularity effect leads to a 19 percent underestimation of target capital in typical small portfolios.

To the best of our knowledge, this is the first paper that develops a unified framework for analyzing a wide range of errors in assessments of portfolio credit risk. Most of the related literature has focused exclusively on misspecifications of the ASRF model and has proposed partial corrections that do not impair the model's tractability. Empirical analyses of violations of the perfect granularity assumption include Martin and Wilde (2002), Vasicek (2002), Emmer and Tasche (2003), and Gordy and Lütkebohmert (2007). For their part, Pykhtin (2004), Düllmann (2006), Düllmann and Masschelein (2006), and Garcia Cespedes et al. (2006) have analyzed implications of the common-factor assumption under different degrees of portfolio concentration in industrial sectors. In addition, Heitfield, Burton, and Chomsisengphet (2006) and Düllman, Scheicher, and Schmieder (2006) have examined both granularity and sector concentration issues in the context of U.S. and European bank portfolios, respectively.[5] A small branch of the related literature, which includes Loeffler (2003) and Morinaga and Shiina (2005), has considered only calibration issues and has derived that noise in model parameters can have a significant impact on assessments of portfolio credit risk.

[5]The recent working paper by the Basel Committee on Bank Supervision (2006) provides an extensive review of these articles.

The remainder of this paper is organized as follows. Section 2 outlines the ASRF model and the empirical methodology applied to it. Section 3 describes the data and section 4 reports the empirical results. Finally, section 5 concludes.

2. Methodology

In this section, we first outline the ASRF model. Then, we discuss how violations of its key assumptions or flawed calibration of its parameters can affect assessments of portfolio credit risk. Finally, we develop an empirical methodology for quantifying and comparing alternative sources of error in such assessments.

2.1 The ASRF Model[6]

The ASRF model of portfolio credit risk—introduced by Vasicek (1991)—postulates that an obligor defaults when its assets fall below some threshold. In addition, the model assumes that asset values are driven by a single common factor:

$$V_{iT} = \rho_i \cdot M_T + \sqrt{1 - \rho_i^2} \cdot Z_{iT}, \tag{1}$$

where V_{iT} is the value of assets of obligor i at time T; M_T and Z_{iT} denote the common and idiosyncratic factors, respectively; and $\rho_i \in [-1, 1]$ is the obligor-specific loading on the common factor. The common and idiosyncratic factors are independent of each other and scaled to random variables with mean 0 and variance 1.[7] Thus, the asset return correlation between borrowers i and j is given by $\rho_i \rho_j$.

The ASRF model delivers a closed-form approximation to the probability distribution of default losses on a portfolio of N exposures. The accuracy of the approximation increases when the number of exposures grows, $N \to \infty$, and the largest exposure weight

[6]This section provides an intuitive discussion of the ASRF model. For a rigorous and detailed study of this model, see Gordy (2003). In addition, Frey and McNeil (2003) examine the calibration of the ASRF model under the so-called Bernoulli mixture representation.

[7]The ASRF model can accommodate distributions with infinite second moments. Nonetheless, we abstract from this generalization in order to streamline the analysis.

shrinks, $\sup_i(w_i) \rightarrow 0$. In these limits, in which the portfolio is perfectly granular, the probability distribution of default losses can be derived as follows. First, let the indicator $\mathcal{I}_{iT}$ equal 1 if obligor i is in default at time T and 0 otherwise. Conditional on the value of the common factor, the expectation of this indicator equals

$$\begin{aligned} E(\mathcal{I}_{iT}|M_T) &= \Pr(V_{iT} < \mathcal{F}^{-1}(PD_{iT})|M_T) \\ &= \Pr(\rho_i \cdot M_T + \sqrt{1-\rho_i^2} \cdot Z_{iT} < \mathcal{F}^{-1}(PD_{iT})|M_T) \\ &= \mathcal{H}\left(\frac{\mathcal{F}^{-1}(PD_{iT}) - \rho_i M_T}{\sqrt{1-\rho_i^2}}\right), \end{aligned}$$

where PD_{iT} is the unconditional probability that obligor i is in default at time T; the cumulative distribution function (CDF) of Z_{iT} is denoted by $\mathcal{H}(\cdot)$; and the CDF of V_{iT} is $\mathcal{F}(\cdot)$, implying that the default threshold equals $\mathcal{F}^{-1}(PD_{iT})$.

Second, under perfect granularity, the Law of Large Numbers implies that the conditional total loss on the portfolio, $TL|M$, is deterministic for any value of the common factor M:

$$\begin{aligned} TL|M &= \sum_i w_i \cdot E(LGD_i) \cdot E(\mathcal{I}_i|M) \\ &= \sum_i w_i \cdot E(LGD_i) \cdot \mathcal{H}\left(\frac{\mathcal{F}^{-1}(PD_i) - \rho_i M}{\sqrt{1-\rho_i^2}}\right), \end{aligned} \tag{2}$$

where time subscripts have been suppressed. In addition, the loss given default of obligor i, LGD_i, is assumed to be independent of both the common and idiosyncratic factors.[8]

Finally, the conditional total loss $TL|M$ is a decreasing function of the common factor M and, consequently, the unconditional distribution of TL can be derived directly on the basis of equation (2) and the CDF of the common factor, $\mathcal{G}(\cdot)$. Denoting by

[8]The ASRF model does allow for interdependence between asset returns and the LGD random variable. Such interdependence leads to another dimension in the study of portfolio credit risk, which is explored by Kupiec (2008). We abstract from this additional dimension in order to focus on the correlation of default events.

$TL_{1-\alpha}$ the $(1-\alpha)^{th}$ percentile in the distribution of total losses, i.e., $\Pr(TL < TL_{1-a}) = 1-\alpha$, it follows that

$$\begin{aligned} TL_{1-\alpha} &= \sum_i w_i \cdot E(LGD_i) \cdot \mathcal{H}\left(\frac{\mathcal{F}^{-1}(PD_i) - \rho_i \mathcal{G}^{-1}(\alpha)}{\sqrt{1-\rho_i^2}}\right) \\ &= TL|M_\alpha, \end{aligned} \tag{3}$$

where $M_\alpha \equiv G^{-1}(\alpha)$ is the α^{th} percentile in the distribution of the common factor. The magnitude $TL_{1-\alpha}$ is also known as the credit VaR at the $(1-\alpha)$ confidence level.

The capital buffer that covers unexpected (i.e., total minus expected) losses on the entire portfolio with probability $(1-\alpha)$ equals[9]

$$\begin{aligned} \kappa &= TL_{1-\alpha} - \sum_i w_i \cdot E(LGD_i) \cdot PD_i \\ &= \sum_i w_i \cdot E(LGD_i) \cdot \left[\mathcal{H}\left(\frac{\mathcal{F}^{-1}(PD_i) - \rho_i \mathcal{G}^{-1}(\alpha)}{\sqrt{1-\rho_i^2}}\right) - PD_i\right] \\ &\equiv \sum_i w_i \cdot \kappa_i. \end{aligned} \tag{4}$$

As implied by this equation, the capital buffer for the portfolio can be set on the basis of exposure-specific parameters, which comprise the exposure's weight in the portfolio, as well as its LGD, PD, and loading on the common factor. The flip side of this implication

[9]In line with the IRB approach of Basel II, this paper effectively incorporates a one-period model, in which a default occurs only at the end of the horizon (in our case, T). In contrast, Gupta et al. (2005) consider obligors that default if their assets are below the default threshold at any one of *multiple* dates over the horizon. We do not consider such a multiperiod variant of the model because one of its key inputs is the term structure of asset return correlations, which cannot be estimated on the basis of our data set (see section 3). Importantly, abstracting from "multiperiod" considerations is unlikely to influence our conclusions. Gupta et al. (2005) find that switching from a one-period to a multiperiod setup lowers a 99.93 percent credit VaR by roughly 5 percent. Although not trivial, this magnitude is several times smaller than the magnitude of errors we identify and analyze below.

is that the portion of the capital buffer attributed to any particular exposure is independent of the rest of the portfolio and, thus, is *portfolio invariant.*[10]

In practice, an implementation of the ASRF model requires that one specify the distribution of the common and idiosyncratic factors of asset returns. It is standard to assume normal distributions and rewrite equation (4) as

$$\kappa = \sum_i w_i \cdot E(LGD_i) \cdot \left[\Phi \left(\frac{\Phi^{-1}(PD_i) - \rho_i \Phi^{-1}(\alpha)}{\sqrt{1-\rho_i^2}} \right) - PD_i \right], \quad (5)$$

where $\Phi(\cdot)$ is the CDF of a standard normal variable. Equation (5) underpins the regulatory capital formula in the IRB approach of Basel II, in which $E(LGD_i) = 45$ percent and $\alpha = 0.1$ percent.

2.2 Impact of Model Misspecification

The portfolio invariance implication of the ASRF model hinges on two key assumptions—i.e., that the portfolio is of perfect granularity and that there is a single common factor. In light of this, we examine at a conceptual level how violations of either of these assumptions—which give rise to "granularity" and "multifactor" effects—affect capital measures. For the illustrative examples in this section, we use the ASRF formula in equation (5) and set $\alpha = 0.1$ percent.

2.2.1 Granularity Effect

The granularity effect arises empirically either because of a limited number of exposures or because of exposure concentration in a small number of borrowers. In either of these cases, idiosyncratic risk is not fully diversified away. Therefore, the existence of a granularity effect implies that capital measures based on the ASRF model would be insufficient to cover unexpected losses.

The top-left panel in figure 1 provides an illustrative example of the granularity effect. This example focuses on a homogeneous portfolio, which gives rise to a desired capital level that decreases with

[10]Given portfolio invariance, the contribution of each exposure to portfolio credit risk is simply $w_i \kappa_i / \kappa$. Ordovas and Thompson (2003) develop a procedure for deriving this contribution in a more general setting.

the number of exposures (solid line).[11] The figure also plots the capital measure obtained under the ASRF model and, thus, incorporates the assumption that there is an infinite number of exposures (dotted line). The difference between the dotted and solid lines equals the magnitude of the granularity effect. As expected, the granularity effect is always negative and decreases when the number of exposures increases.

Gordy and Lütkebohmert (2007) derive a closed-form "granularity adjustment," which approximates (the negative of) the granularity effect. When the portfolio is homogeneous, the approximation is linear in the reciprocal of the number of exposures, which is largely in line with the properties of the granularity effect plotted in figure 1.[12]

2.2.2 Multifactor Effect

The impact of various macroeconomic and industry-specific conditions on portfolio credit risk may be best accounted for by generalizing equation (1) to incorporate *multiple* (potentially unobservable) common factors. Multiple common factors affect the likelihood of default clustering (i.e., the likelihood of a large number of defaults occurring over a given horizon), which influences the tails of the probability distribution of credit losses.[13] In line with our empirical results (reported in section 4 below), we treat a fattening of these tails (reflected in greater kurtosis) as implying unambiguously a higher level of the desired capital buffer.[14]

Importantly, the existence of multiple common factors of credit risk would violate the single-factor assumption of the ASRF model, leading to what we call a multifactor effect in capital measures. Depending on the characteristics of the credit portfolio, the multifactor effect could be either negative—i.e., implying that the ASRF

[11] The calculation of the desired capital level uses a Gaussian copula (see appendix 1).

[12] Further comparison between the granularity effect and the granularity adjustment of Gordy and Lütkebohmert (2007) is reported in section 4.1.

[13] These losses were denoted by TL in section 2.1.

[14] Note that, depending on the confidence level of the targeted credit VaR, a fattening of the tails of the probability distribution of credit losses may either raise or lower the desired capital buffer. Fatter tails of the loss distribution translate into a higher desired level of the capital buffer only if the value of α in equation (4) is sufficiently close to zero.

Figure 1. Four Sources of Error in Capital Measures

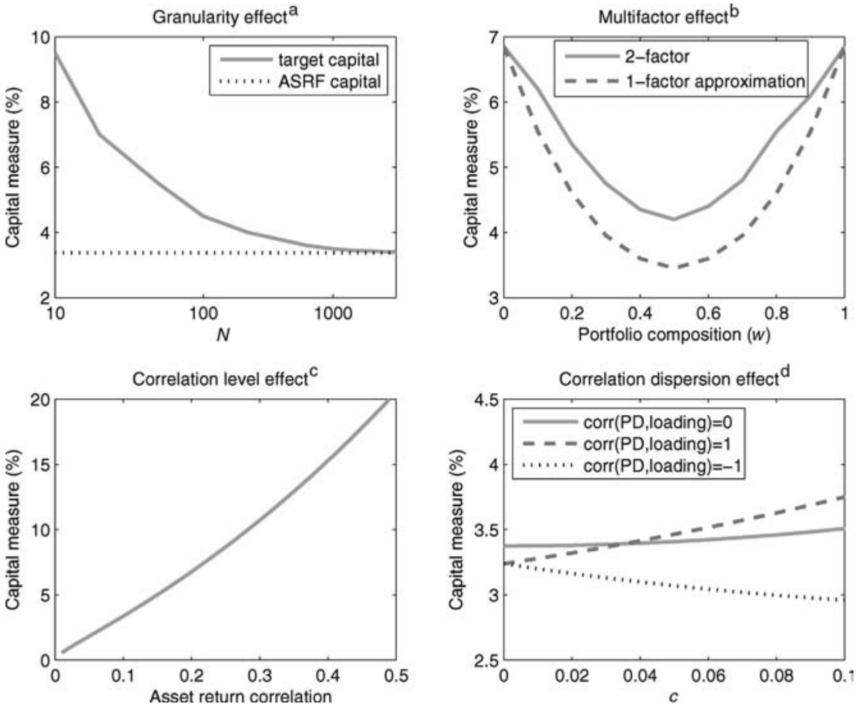

Note: Capital measures, in percent and per unit of aggregate exposure, are shown on the vertical axes. For each panel (unless noted otherwise), PD = 1 percent and LGD = 50 percent are the same across exposures.
[a]The solid line plots target capital for a portfolio in which all pairwise asset return correlations equal 10 percent. The number of exposures in the portfolio (N) varies across the horizontal axis. The dotted line plots the corresponding capital estimate when $N \to \infty$. [b]The portfolio consists of 1,000 exposures that are divided into two groups, with w denoting the weight of the first group. Within each group, the asset return correlations equals 20 percent for all exposure pairs. Intergroup correlations are zero. The solid line plots the target capital level, which incorporates the two common factors in the simulated data. The dashed line plots the capital calculated under a one-common-factor approximation of the correlation structure (see appendix 3). [c]Capital measures are shown for different levels of homogeneous pairwise asset return correlations. [d]The solid line plots capital measures under the assumption that PD = 1 percent, there is a single common factor, and the loadings on this factor are distributed uniformly in the cross-section between $\sqrt{0.1} - c$ and $\sqrt{0.1} + c$. For the other two lines, PDs are distributed uniformly in the cross-section between 0.5 percent and 1.5 percent and have a positive (dashed line) or negative (dotted line) linear relationship with the common-factor loadings.

model underestimates the desired capital—or positive. To illustrate the two possibilities, we generalize equation (1) to

$$V_i = \rho_{1,i} \cdot M_1 + \rho_{2,i} \cdot M_2 + \sqrt{1 - \rho_{1,i}^2 - \rho_{2,i}^2} \cdot Z_i, \tag{6}$$

where M_1, M_2, and Z are mutually independent standard normal variables.

In our first example, we consider a portfolio in which all exposures have equal weights, have the same PD, and are divided into two groups according to their dependence on the common factors. For exposures in the first group, $0 < \rho_{1,i} = \rho < 1$ and $\rho_{2,i} = 0$, while $\rho_{1,j} = 0$ and $0 < \rho_{2,j} = \rho < 1$ for exposures in the second group. Thus, the common factors are group specific and underpin positive and homogeneous within-group pairwise correlations and zero across-group correlations. The solid line in the top-right panel of figure 1 plots the desired capital measure for such a portfolio as a function of the relative weight of exposures in group 1.[15] This measure is lowest when the portfolio is most diversified between the two groups of exposures and, thus, the probability of large default losses is minimized. In addition, the dashed line in the panel plots an alternative capital measure, which is based on the ASRF model and, thus, incorporates a single-factor structure of the asset return correlations. We choose this structure optimally in order to provide the ASRF model with as much information about the true correlations as possible, subject to the single-factor constraint (see appendix 3). As it turns out, the single-factor approximation matches extremely well the true *average* asset return correlation but approximates only roughly the dispersion of correlation coefficients in the cross-section of exposures.

[15]The desired capital buffer is calculated on the basis of Monte Carlo simulations (see appendix 2) for a portfolio consisting of 1,000 exposures.

The difference between the dashed and solid lines equals the multifactor effect. This effect is *negative* because the single-factor assumption of the ASRF model ignores the fact that the common factors are two independent sources of default clustering, which leads to an underestimation of the desired capital. The underestimation is largest when the two groups enter the portfolio with equal weights, in which case the role of multiple factors is greatest.

It is possible, however, to construct another example, in which imposing an erroneous single-factor structure on portfolio credit risk distorts the interaction between asset return correlations and individual PDs in a way that leads to a *positive* multifactor effect. Consider a portfolio comprising two groups of exposures, with the exposures in the first group being individually riskier but less correlated among themselves than the exposures in the second group. In terms of equation (6), this can be formalized by postulating that firms with high PDs feature $0 < \rho_{1,i} = \rho < 1$ and $\rho_{2,i} = 0$, whereas firms with low PDs feature $0 < \rho_{1,j} = \rho_{2,j} = \rho < 1$. A single-factor approximation to this correlation structure would match the average correlation coefficient but would also imply too high a correlation among riskier exposures. Raising the probability of default clustering, this would lead to a capital buffer that is larger than desired.

2.3 Impact of Calibration Errors

Errors in the calibration of the ASRF model will affect assessments of portfolio credit risk even if this model is well specified. In this paper, we focus on errors in the calibration of the interdependence of credit risk across exposures, which can be driven by noise in the adopted values of asset return correlations or by a flawed assumption regarding the distribution of asset returns. When analyzing the consequences of such errors, we maintain our earlier practice and treat fattening of the tails of the loss distribution as implying unambiguously a higher 99.9 percent credit VaR and, thus, a higher desired level of the capital buffer. In this way, we sharpen the conceptual analysis and keep it in line with our empirical findings.

We study two general types of errors in calibrated asset return correlations: errors in the average correlation coefficient and errors in the dispersion of correlation coefficients across exposure pairs. It is important to keep in mind how these calibration-driven errors relate to errors induced by the two misspecification effects that we examined in sections 2.2.1 and 2.2.2. First, given that they occur within the constraints of the ASRF model, calibration errors are independent of the granularity effect. Second, given the properties of the one-factor approximation that underlies the extraction of the multifactor effect (see section 2.2.2), this effect is effectively separated from errors in the average correlation. By contrast, errors in the calibrated dispersion of asset return correlations could arise either as a result of ignoring the importance of multiple factors (recall section 2.2.2) or as a result of noise in the estimated factor loadings when there is a single common factor. In this section, we are concerned with the second case, as it is consistent with a correct specification of the ASRF model and refers only to calibration errors.

The two types of errors in calibrated asset return correlations have various potential sources. One possibility is that a user of the ASRF model relies entirely on rule-of-thumb values, which may simply be correlation estimates for a popular credit index. Such estimates will lead to a discrepancy between desired and calculated capital to the extent that the popular index is not representative of the user's own portfolio. Alternatively, a user of the model may have access only to short time series of data on the assets of the obligors in its portfolio, which would lead to small-sample estimation errors in asset return correlations. Indeed, this second source of error is likely to be important in practice because (i) asset value estimates are typically available at low (i.e., monthly or quarterly) frequencies and (ii) supervisory texts require that financial institutions possess only five years of relevant data.[16]

A positive error in the average level of asset return correlations leads to a capital measure that is higher than the desired one (figure 1, bottom-left panel). This result reflects the intuition that inflating asset return correlations increases the likelihood of default clustering, which fattens the tails of the loss distribution. In the remainder

[16]Data limitations are likely to be important irrespective of how a user of the model estimates asset return correlations. Such estimates may rely on balance sheet information and stock market data. Alternatively, as derived in Tarashev and Zhu (2006), asset return correlations can be extracted from the CDS market.

of this paper, the impact of errors in the *average* correlation on capital measures is dubbed the "correlation level" effect.

In turn, the effect of noise in the estimated dispersion of correlation coefficients can be seen in the following 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 characterized 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, in which pairwise correlations exhibit dispersion, is more likely to experience several simultaneous defaults and, thus, has a loss distribution with fatter tails. Consequently, of the two portfolios, the second one requires higher capital in order to attain solvency with the same probability. This is portrayed by the upward slope of the solid line in the bottom-right panel of figure 1 and is a particular instance of what we dub the "correlation dispersion" effect, which arises in the context of a single common factor.

This result can be strengthened (dashed line in the same panel) but also weakened or even reversed if PDs vary across firms. To see why, consider the previous example but 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. As a result, greater dispersion of asset return correlations may lower the probability of default clustering in the second portfolio to an extent that depresses the desired capital level below that for the first portfolio. This is illustrated by the negative slope of the dotted line in the bottom-right panel of figure 1.

Even if asset return correlations were known with certainty, a flawed calibration of the overall distribution of asset returns would still drive errors in the calibrated interdependence of credit risk across exposures. Although the ASRF model imposes quite weak restrictions on asset return distributions, it is common practice to adopt distributions whose main advantages stem not from realistic features but from operational convenience. In particular, the consensus in the literature is that asset returns have fatter tails than those imposed by the conventional normality assumption. To the extent that the fatness of the tails reflects the distribution of the common factor, the probability of default clustering and, thus, the desired capital level would be higher than those derived under normality (Hull and White 2004; Tarashev and Zhu 2006). We study this issue by considering Student-t distributions for both the common and idiosyncratic factors of asset returns.

2.4 Evaluating Various Sources of Error

An important contribution of this paper is to present a unified empirical method for quantifying the impact of several sources of error in model-based assessment of portfolio credit risk. In particular, we focus on the difference between target capital measures and shortcut ones, the latter of which are based on the ASRF model and probable erroneous calibration of its parameters. We dissect this difference into four non-overlapping and exhaustive components, attributing them to the multifactor, granularity, correlation level, and correlation dispersion effects. In order to probe further the likely magnitude of the last two effects, we derive plausible small-sample errors that could affect direct estimates of asset return correlations. Finally, we also examine the implications of erroneous assumptions regarding the distribution of asset returns.

2.4.1 Baseline Method

The baseline empirical method consists of two general steps. In the first step, we construct a hypothetical portfolio that is either "large"—consisting of 1,000 equal exposures—or "small"—consisting of 200 equal exposures.[17] The sectoral composition of such a portfolio is constrained to be in line with the typical loan portfolio of large wholesale banks in the United States.[18] Given the constraint, the portfolio is sampled at random from the entire population of firms in our data set. Since each simulated portfolio is subject to sampling noise, we examine 3,000 different draws for both large and small portfolios.

[17] The distinction between what we dub large and small portfolios does not reflect the size of the aggregate exposure but rather different degrees of diversification across individual exposures (see section 3.2 for further detail).

[18] Such a portfolio does not incorporate consumer loans and, thus, may not be representative of all aspects of credit risk.

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. Each of these alternatives employs the same set of PD values and assumes that asset returns are normally distributed. In addition, each alternative incorporates random LGDs that are independent of all random variables driving defaults. The LGD distributions are symmetric, triangular, peak at 50 percent, have a continuous support on the interval $[0\%, 100\%]$, and are identical and independent across exposures.[19]

Each measure differs from a previous one owing to a single assumption:

1. The *target capital* measure incorporates data on asset return correlations, which are treated as representing the "truth." Using these correlations, we conduct Monte Carlo simulations to construct the "true" probability distribution of default losses at the one-year horizon. The implied 99.9 percent credit VaR minus expected losses equals target capital (see appendix 2 for further detail).[20]
2. The second capital measure differs from target capital only owing to a restriction on the number of common factors governing asset returns. In particular, we adopt a correlation matrix that fits the original one as closely as possible under the constraint that correlation coefficients should be consistent with the presence of a single common factor (see appendix 3). The fitted single-factor correlation matrix is used to derive the one-year probability distribution of joint defaults

[19] The LGD specification warrants an explanation. The independence between the incidence of defaults and LGDs implies that, in the absence of simulation noise, only the mean of the LGD distribution enters (as a multiplicative factor) capital measures. However, the entire LGD distribution affects measures obtained from Monte Carlo simulations. Importantly, assuming a continuous distribution for LGDs smoothes the derived probability distribution of joint defaults, which improves the robustness of simulation-based capital measures.

[20] In order to reduce simulation errors to levels that do not affect our conclusions, we estimate each portfolio loss distribution on the basis of 500,000 draws. Alternatively, it is possible to circumvent the use of Monte Carlo simulations by approximating loss distributions analytically, on the basis of the so-called saddle-point method of Browne, Martin, and Thompson (2001).

on the basis of the so-called Gaussian copula method (see appendix 1). This distribution is then mapped into a probability distribution of default losses and, finally, into a capital measure.

3. The third capital measure differs from the second one only in that it assumes that all idiosyncratic risk is diversified away. This assumption allows us to use the fitted single-factor correlation matrix, which underpins measure 2 in the ASRF formula (equation (5)).
4. The fourth capital measure differs from the third one only in that it is based on the assumption that loading coefficients on the single common factor are the same across exposures. The resulting common correlation coefficient, which is set equal to the average of the pairwise correlations underpinning measures 2 and 3, is used as an input to the ASRF formula (equation (5)).
5. Finally, the *shortcut* capital measure differs from the fourth one only in that it incorporates alternative, rule-of-thumb, values for the common correlation coefficient.

The three intermediate measures lead to a straightforward dissection of the difference between target and shortcut capital.[21] Specifically, the difference between measures 5 and 1 is the sum of the following four components: (i) the difference between measures 2 and 1, which equals the multifactor effect; (ii) the difference between measures 3 and 2, which equals the granularity effect; (iii) the difference between measures 4 and 3, which equals the correlation dispersion effect; and (iv) the difference between measures 5 and 4, which equals the correlation level effect.

The specific ordering and choice of the three intermediate capital measures is a result of the following reasoning. As far as *measure 2* is concerned, it is determined by the necessity to extract the multifactor effect first. The reason for this is twofold. First, deriving a capital measure that assumes an infinite number of exposures but allows for

[21] Importantly, the method also applies to alternative definitions of *target* and *shortcut* capital, so long as the *true* correlation structure and shortcut correlation estimates chosen by the user are well defined.

multiple factors (i.e., extracting the granularity effect before the multifactor effect) is subject to approximation errors (see, e.g., Pykhtin 2004). Second, it is possible to isolate calibration errors (via measures 4 and 5) only after the extraction of the multifactor effect (via measure 2) has modified the original correlation matrix so that it is consistent with the ASRF model. Likewise, an application of this model for the extraction of calibration errors requires the assumption of infinite granularity, which explains why measure 3 is calculated before measures 4 and 5.[22] Finally, modifying measure 4 by preserving the cross-sectional distribution of the single-factor correlation coefficients but changing their average level would reverse the order in which the correlation level and dispersion effects are extracted. An important problem with this procedure is that it would allow only for an imperfect estimate of the correlation level effect because this estimate would be influenced by changes in the *structure* of the single-factor correlation matrix.

2.4.2 Two Extensions

In an attempt to delve further into the impact of *plausible* calibration errors on capital measures, we conduct two additional exercises. Each exercise focuses on a specific type of error in the calibrated interdependence of defaults and incorporates the assumption that the true PDs are identical across exposures. This assumption insulates capital measures from the impact of interaction between heterogeneous PDs and errors in the calibrated interdependence of defaults.

In the first exercise, we derive the extent to which plausible limitations on the size of available data can affect assessments of portfolio credit risk by affecting the estimates of asset return correlations. Specifically, we draw time series of asset returns from a joint distribution in which all pairwise correlations equal the correlation underpinning measure 4. Then we use the sample correlation matrix of the simulated series, a typical value for the probability of default,

[22] Despite this observation, we have also experimented with imposing the infinite granularity assumption only after we calculate the correlation dispersion and level effects (i.e., deriving measure 3 after measures 4 and 5). This modifies the meaning of these two effects but alters negligibly their magnitudes, as well as the magnitude of the granularity effect.

and the ASRF formula in equation (5) in order to derive a capital measure.[23] The difference between this measure and the desired capital, which employs the *exact* correlation structure, is driven by small-sample noise in the estimates of asset return correlations.

In our second exercise, we examine how measure 4 would change if the common and idiosyncratic factors of asset returns are in fact driven by Student-t distributions. The results of this exercise reveal how flawed calibration of the tail-of-distribution dependence among asset returns affects capital calculations. In order to carry out the exercise, we use the general ASRF formula in equation (4) and make two technical adjustments to the empirical setup. The first adjustment rescales the common and idiosyncratic factors to random variables with a unit variance.[24] The second adjustment addresses the fact that the generalized CDF of asset returns, $\mathcal{F}(\cdot)$, does not exist in closed form. In concrete terms, we calculate the default threshold $\mathcal{F}^{-1}(PD)$ on the basis of 10 million Monte Carlo simulations.

3. Data Description

This section describes the two major blocks of data that we rely on: (i) credit-risk parameter estimates provided by Moody's KMV and (ii) the sectoral distribution of exposures in typical portfolios of U.S. wholesale banks.

3.1 Credit-Risk Parameters

Our sample includes the universe of firms covered in July 2006 by both the expected default frequency (EDF$^{\text{TM}}$) model and the global correlation (GCorr$^{\text{TM}}$) model of Moody's KMV. These two models deliver, respectively, estimates of one-year physical PDs and physical asset return correlation coefficients for publicly traded companies. We abstract from financial firms and work with 10,891 companies.

The sample covers firms with diverse characteristics. Specifically, 5,709 of the firms are headquartered in the United States; 4,383 in Western Europe; and the remaining 799 in the rest of the world. Further, the distribution of the 10,891 firms across

[23] This measure abstracts from the granularity and multifactor effects.

[24] A Student-t variable with $r > 2$ degrees of freedom has a variance of $\frac{r}{r-2}$.

industrial sectors is reported in the last column in table 1, with the largest share of firms (10.4 percent) coming from the business service sector. Importantly, only 1,434 (or 13.2 percent) of the firms have a rating from either S&P or Moody's, which matches the stylized fact that the majority of bank exposures are unrated.

There are several reasons why EDFs and GCorr correlations are natural data for our exercise. First, the two measures are derived within mutually consistent frameworks, which build on the model of Merton (1974) and are in the spirit of the ASRF model (see Das and Ishii 2001, Crosbie and Bohn 2003, and Crosbie 2005 for detail). Second, in line with their role in this paper, Moody's KMV EDFs have been widely used as proxies for actual default probabilities (see, e.g., Berndt et al. 2005 and Longstaff, Mithal, and Neis 2005). Third, the GCorr correlations are underpinned by a multifactor structure, which is crucial for our study of the multifactor effect. In particular, this model incorporates 120 common factors that comprise 2 global economic factors, 5 regional economic factors, 7 sector factors, 61 industry-specific factors, and 45 country-specific factors.

Table 2 and figure 2 report summary statistics of the Moody's KMV one-year PD and asset return correlation estimates. The cross-sectional distribution of PDs has a long right tail and, thus, its median (0.39 percent) is much lower than its mean (2.67 percent). In addition, the favorable credit conditions in July 2006 resulted in 1,217 firms (i.e., about 11.2 percent of the total) having the lowest EDF score (0.02 percent) allowed by the Moody's KMV empirical methodology. At the same time, the upper bound on the Moody's KMV PD estimates (20 percent) is attained by 643 firms. For their part, GCorr correlations are limited between 0 and 65 percent. Clustered mainly between 5 percent and 25 percent, these correlations average 9.24 percent.[25]

[25]The GCorr correlation estimates are quite in line with correlation estimates reported in other studies. For instance, Lopez (2004) documents an average asset correlation of 12.5 percent for a large number of U.S. firms and Düllman, Scheicher, and Schmieder (2006) estimate a median asset return correlation of 10.1 percent for European firms.

Table 1. Sectoral Composition of Simulated Portfolios

	Large Portfolio		Small Portfolio		
Sector	**Number of Names**	**Exposure Weight (%)**	**Number of Names**	**Exposure Weight (%)**	**Number of Firms in the Sample**
Aerospace and Defense	31	3.1			105
Agriculture	9	0.9			56
Air Transportation	4	0.4			83
Apparel, Footwear, and Textiles	17	1.7			357
Automotive	51	5.1	19	9.5	198
Broadcast Media	43	4.3	16	8.0	191
Business Services	23	2.3			1,132
Chemicals	42	4.2	15	7.5	940
Computer Equipment	10	1.0			746
Construction	43	4.3	16	8.0	277
Electric, Gas, and Sanitary	107	10.7	39	19.5	335
Electronics and Electrical	18	1.8			693
Entertainment and Leisure	33	3.3			294
Fabricated Metals	17	1.7			146
Food, Beverages, and Tobacco	63	6.3	23	11.5	490
General Retail	31	3.1			133
Glass and Stone	6	0.6			149
Health Care	38	3.8	14	7.0	178
Legal and Other Services	16	1.6			452
Lodging	22	2.2			70
Machinery and Equipment	36	3.6	13	6.5	645

(continued)

Table 1. (Continued)

Sector	Large Portfolio		Small Portfolio		
	Number of Names	**Exposure Weight (%)**	**Number of Names**	**Exposure Weight (%)**	**Number of Firms in the Sample**
Medical Equipment	10	1.0			334
Mining	6	0.6			486
Miscellaneous Manufacturing	18	1.8			130
Nondefense Trans. and Parts	2	0.2			53
Oil and Gas Exploration	55	5.5	20	10.0	100
Oil Refining and Delivery	24	2.4			458
Other Trans. Services	22	2.2			108
Paper and Forestry	23	2.3			172
Personal Services	7	0.7			31
Primary Metals	11	1.1			188
Printing and Publishing	28	2.8			186
Repair Services and Rental	13	1.3			37
Restaurants	9	0.9			141
Rubber and Plastics	18	1.8			120
Semiconductors	2	0.2			177
Telecommunications	69	6.9	25	12.5	177
Trucking and Warehousing	5	0.5			68
Water Transportation	5	0.5			104
Wood, Furniture, and Fixtures	13	1.3			151
Total	1,000	100	200	100	10,891
Sector Concentration Index	0.0432		0.1135		
Name Concentration Index	0.0010		0.0050		

Note: A concentration index equals the sum of squared weights, set either at the firm or sector level.

Table 2. Summary Statistics (in percent)

	Mean	Std. Dev.	Skewness	Median	Minimum	Maximum
One-Year PDs	2.67	5.28	2.49	0.39	0.02	20.00
Pairwise Correlations	9.24	3.86	1.87	8.45	0.29	65.00

Note: The sample includes 10,891 nonfinancial firms.

Figure 2. Distribution of Individual PDs and Pairwise Correlations

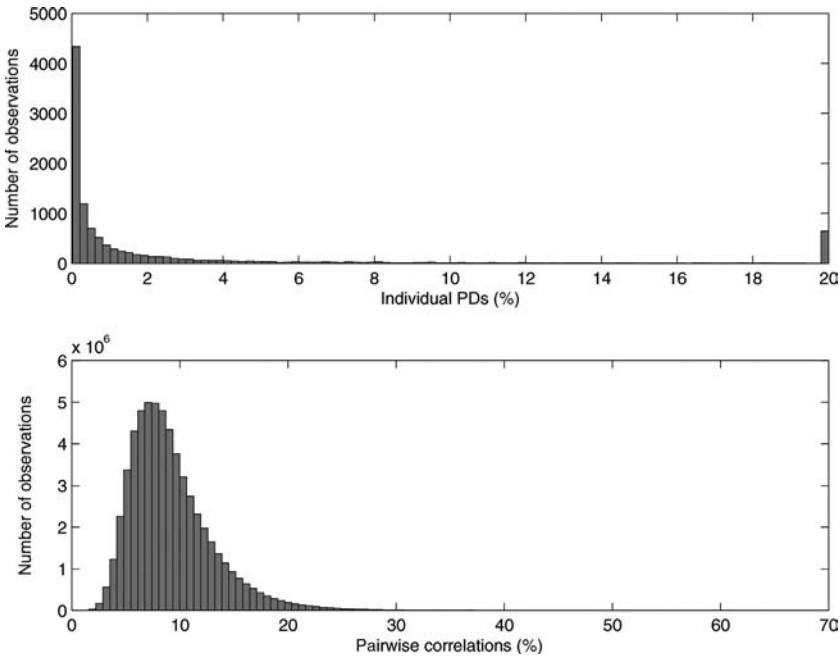

Note: The parameter estimates relate to 10,891 nonfinancial firms in July 2006.
Data Source: Moody's KMV.

3.2 Characteristics of Hypothetical Portfolios

The portfolios we simulate match the sectoral distribution of the typical portfolio of U.S. wholesale banks. Specifically, to construct a large portfolio (1,000 exposures), we apply the forty nonfinancial sector weights reported by Heitfield, Burton, and Chomsisengphet (2006) (see table 1). For a small portfolio (200 exposures), we rescale the ten largest sectoral weights so that they sum to unity, and we set all other weights to zero. Within each sector, we draw firms at random.[26] All firms in a portfolio receive equal weights (up to a rounding error) and, thus, there is a one-to-one correspondence between the number of firms in a sector and that sector's weight in the portfolio.

The sector and name concentration indices, reported at the bottom of table 1, provide justification for our design of large and small portfolios. Calculated as the sum of squared sectoral (name) weights, the sector (name) concentration index of the large portfolios studied in this paper equals 0.0432 (0.001). This belongs to the range of concentration indices $[0.03, 0.045]$ ($[0.000, 0.003]$) reported in Heitfield, Burton, and Chomsisengphet (2006) for large portfolios of U.S. banks. For small portfolios, the analogous indices and ranges are 0.1135 (0.005) and $[0.035, 0.213]$ ($[0.001, 0.008]$), respectively.

4. Empirical Results

We implement the empirical methodology described in section 2.4 in order to quantify the impact of various sources of error in ASRF-based assessments of portfolio credit risk. Before reporting our findings, it is useful to highlight several aspects of the methodology.

First, as far as calibration of the model is concerned, the analysis in this paper focuses exclusively on errors in the values of parameters that relate to the interdependence of credit risk across exposures. Considering the impact of noise in PD and LGD estimates would make it extremely difficult to isolate the correlation level and dispersion effects we focus on. This is because noise in PDs and LGDs would interact with noise in correlation inputs in a highly nonlinear fashion.

[26] Within each industry sector, we draw randomly *with* replacement. If the same firm is drawn twice, the corresponding pairwise correlation is set equal to the average correlation for the sector. Drawing randomly *without* replacement does not affect materially the results.

Second, we make the stylized assumption that portfolios consist of equally weighted exposures. Considering disparate exposure sizes would require considering an additional dimension of portfolio characteristics, as it will no longer be the case that the granularity of a larger portfolio is necessarily finer. In addition, lower granularity that results from higher concentration in a small number of borrowers would also have a bearing on the number and importance of common factors affecting the portfolio and on the overall correlation of risk across exposures. This would make it impossible to isolate the granularity effect from the other three effects we consider.

Finally, our analysis treats the correlation matrix provided by Moody's KMV as revealing the "true" correlation of asset returns. Of course, this matrix is itself an estimate that is subject to errors. Nevertheless, the Moody's KMV correlation matrix provides a reasonable benchmark to work from. In addition, we have verified that results regarding the *relative* importance of alternative sources of error depend only marginally on the accuracy of the GCorr estimates, even though the *absolute* impact of alternative sources of error does change with the benchmark correlation level.

4.1 Various Errors in Shortcut Capital Measures

To study various sources of error in assessments of portfolio credit risk, we calculate the five capital measures listed in section 2.4 for 3,000 large and as many small hypothetical portfolios. Even though they have the same sectoral composition by construction, the simulated portfolios differ from each other with respect to the individual constituent exposures and, thus, with respect to the underlying risk parameters (see table 3).

Table 4 reports summary statistics of the target and shortcut capital measures (i.e., the two extremes described in section 2.4). For large portfolios, the target capital averages 3.31 percent (per unit of aggregate exposure) across the 3,000 simulated portfolios. The corresponding shortcut level (based on a rule-of-thumb asset return correlation of 12 percent) is 81 basis points higher.

Decomposing the difference between target and shortcut capital for large portfolios reveals that errors caused by model misspecification play a minor role. In qualitative terms, the multifactor effect

Table 3. Characteristics of Simulated Loan Portfolios (in percent)

	Mean	Std. Dev.	Median	Minimum	Maximum
A. Large Portfolios (1,000 Firms)					
Average PD	2.42	0.19	2.42	1.79	3.12
Std. Dev. of Individual 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	9.78	0.22	9.77	9.14	10.73
Std. Dev. of Loadings	9.33	0.31	9.32	8.33	10.47
Corr. (PD, Loadings)	−20.00	2.04	−20.10	−26.70	−12.80
B. Small Portfolios (200 Firms)					
Average PD	2.28	0.36	2.26	1.24	3.68
Std. Dev. of Individual 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	10.49	0.44	10.48	8.99	12.00
Std. Dev. of Loadings	10.54	0.70	10.55	7.80	12.79
Corr. (PD, Loadings)	−19.80	4.59	−20.20	−31.80	−1.20

Note: The results are based on 3,000 simulated portfolios and are obtained in two steps. First, portfolio-specific characteristics (specified by row headings) are calculated for each simulated portfolio. Second, summary statistics (specified by column headings) are calculated for each of the characteristics obtained in the first step. "Loadings" are estimated under a one-common-factor approximation of the correlation structure and refer to the firm-specific loadings of asset returns on the single common factor.

Table 4. Capital Measures and Four Sources of Errors (in percent)

A. Large Portfolios (1,000 Firms)					
	Mean	**Std. Dev.**	**Median**	**95% Interval**	**50% Interval**
Target Capital[a]	**3.31**	**0.18**	**3.30**	**[2.97, 3.67]**	**[3.19, 3.42]**
Deviation from Target Due To:[b]					
Multifactor Effect[c]	−0.03	0.02	−0.03	[−0.08, 0.01]	[−0.05, −0.02]
Granularity Effect[d]	−0.16	0.004	−0.15	[−0.16, −0.15]	[−0.16, −0.15]
Correlation Dispersion Effect[e]	0.39	0.05	0.39	[0.30, 0.48]	[0.36, 0.42]
Correlation Level Effect[f]	0.61	0.06	0.62	[0.49, 0.73]	[0.57, 0.66]
Shortcut Capital (corr = 12%)	**4.12**	**0.20**	**4.12**	**[3.75, 4.51]**	**[3.99, 4.25]**
Correlation Level Effect if:					
corr = 6%	−1.06	0.08	−1.06	[−1.23, −0.92]	[−1.11, −1.01]
corr = 18%	2.24	0.10	2.24	[2.05, 2.43]	[2.17, 2.30]
corr = 24%	3.86	0.14	3.86	[3.59, 4.14]	[3.77, 3.95]
B. Small Portfolios (200 Firms)					
	Mean	**Std. Dev.**	**Median**	**95% Interval**	**50% Interval**
Target Capital[a]	**3.86**	**0.32**	**3.85**	**[3.25, 4.55]**	**[3.64, 4.07]**
Deviation from Target Due To:[b]					
Multifactor Effect[c]	−0.04	0.04	−0.04	[−0.12, 0.04]	[−0.07, −0.01]
Granularity Effect[d]	−0.73	0.03	−0.73	[−0.79, −0.67]	[−0.75, −0.71]
Correlation Dispersion Effect[e]	0.42	0.12	0.42	[0.19, 0.65]	[0.34, 0.50]
Correlation Level Effect[f]	0.40	0.13	0.40	[0.16, 0.67]	[0.32, 0.48]
Shortcut Capital (corr = 12%)	**3.91**	**0.38**	**3.90**	**[3.17, 4.70]**	**[3.64, 4.16]**

(continued)

Table 4. (Continued)

B. Small Portfolios (200 Firms)					
	Mean	**Std. Dev.**	**Median**	**95% Interval**	**50% Interval**
Correlation Level Effect if:					
corr = 6%	−1.19	0.13	−1.19	[−1.46, −0.95]	[−1.28, −1.10]
corr = 18%	1.95	0.21	1.94	[1.57, 2.38]	[1.81, 2.08]
corr = 24%	3.50	0.30	3.49	[2.94, 4.12]	[3.30, 3.70]

Note: Summary statistics are for the simulated portfolios underpinning table 3. 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. In all calculations, LGD_i has a symmetric triangular distribution between 0 and 1.

[a]Target capital is 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.

[b]Four sources of deviation from the target capital level are shown; 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.

[c]For the multifactor effect, the correlation matrix underpinning the target capital level is approximated under the assumption that there is a single common factor.

[d]For the granularity effect, there is the additional assumption that the number of firms is infinite.

[e]For the correlation dispersion effect, the additional assumption is that the loadings on the single common factor are the same across exposures.

[f]For the correlation level effect, the additional assumption imposes a different, shortcut level on the constant pairwise correlation.

can be of either sign (fourth column in table 4) but is more likely to be negative (fourth and fifth columns in table 4). In light of the discussion in section 2.2.2, imposing a single-factor framework is more likely to lead to too low a capital buffer because such a framework ignores the existence of multiple sources of default clustering. In quantitative terms, however, the multifactor effect entails an average discrepancy that amounts to less than 1 percent of the average target capital level.[27] This is because the single-factor approximation fits closely the raw correlation matrix. Indeed, our single-factor approximation matches almost perfectly the level of average correlations (with a maximum discrepancy across simulated portfolios of less than 4 basis points) and explains on average 76 percent of the variability of pairwise correlations in the cross-section of exposures.[28,29]

Similarly, the granularity effect is with the expected negative sign but, for large portfolios, leads to a small deviation from target capital. With an average of −16 basis points (or roughly 5 percent of target capital), this deviation is nonetheless significantly higher than that induced by the multifactor effect. Not surprisingly, the granularity effect we calculate is approximated extremely well by (the negative of) the closed-form granularity adjustment of Gordy and Lütkebohmert (2007), which averages −17 basis points for large portfolios in our sample. In addition, the correlation between the granularity effect and the granularity adjustment across large simulated portfolios is 66 percent.

By contrast, erroneous calibration of the ASRF model leads to much greater deviations from target. For large portfolios, the correlation dispersion effect raises the capital measure by 39 basis points,

[27]A similar result is obtained by Düllmann and Masschelein (2007), who rely on Pykhtin (2004) to approximate the multifactor effect in loan portfolios of German banks.

[28]The goodness-of-fit measure for the one-factor approximation is described in appendix 3. Across the 3,000 simulations of large portfolios, this measure ranges between 67 and 85 percent. For small portfolios, the range is between 63 and 86 percent.

[29]Principal-component analysis confirms this result. Specifically, the portion of the total variance of asset returns explained by the first principal component is at least ten times larger than the portion explained by the second principal component.

which amounts to roughly 12 percent of the target level.[30] The sign of this effect reflects the fact that, in our data, exposures with higher PDs tend to be less correlated with the rest of the portfolio (see the last row in each panel of table 3).[31] The shortcut capital measure ignores this regularity and, in line with the intuition provided in section 2.3, overshoots the target.

The correlation level effect has a similarly important implication. Specifically, this effect reveals that raising the average correlation coefficient from 9.78 percent (the one observed in the data) to a rule-of-thumb value of 12 percent leads to an 18 percent overestimation of the target capital level. The sign of the deviation is not surprising in light of the discussion in section 2.3. Importantly, the shortcut measure drops (rises) by roughly 8 percent with each percentage point decrease (increase) in the homogeneous correlation coefficient. Thus, using a rule-of-thumb correlation of 6 percent leads to a 32 percent *underestimation* of the target level.[32]

Turning to small portfolios, the decomposition results are qualitatively the same, with the notable exception of the granularity effect. In these portfolios, a much smaller portion of the idiosyncratic risk is diversified away and the granularity effect equals −73 basis points, which implies a 19 percent underestimation of the target capital. This underestimation is approximated well by the Gordy and Lütkebohmert (2007) granularity adjustment, (the negative of) which averages −86 basis points for small portfolios and exhibits an 89 percent correlation with the corresponding granularity effect.

4.2 Regression Analysis of Calibration Errors

Given the dominant role of correlation level and dispersion effects as determinants of model-based assessments of portfolio credit risk, we investigate the sources of these two effects via a regression analysis.

[30]On the basis of a hypothetical portfolio of U.S. firms, Hanson, Pesaran, and Schuermann (2007) also demonstrate the importance of accounting for cross-sectional heterogeneity in credit-risk parameters.

[31]The negative relationship between PDs and correlations (or loading coefficients in a single-factor setting) is likely to be a general phenomenon. See, e.g., Lopez (2004), Arora, Bohn, and Korablev (2005), and Dev (2006), who find that global factors often play bigger roles for firms of better credit quality.

[32]The rule-of-thumb asset return correlations reported in the literature range between 5 and 25 percent.

The regressions—run on the cross-section of simulated portfolios—are simple linear models of *calibration-driven capital discrepancy*, which is defined as shortcut capital (based on a correlation of 12 percent) minus target capital net of the multifactor and granularity effects.

We consider two blocks of explanatory variables. The first block comprises the average level and the dispersion of the asset return correlation coefficients underlying each simulated portfolio.[33] These variables are natural drivers of the correlation level and dispersion effects and would explain the two effects completely if assessments of portfolio credit risk did not depend on the interaction of asset return correlations with PDs. In order to account for such interaction, we include a second block of explanatory variables, which comprises average PDs and the cross-sectional correlations between PDs and single-factor loading coefficients. One would recall that the PDs underlying target capital are identical to those underlying the shortcut measure. Thus, the regression coefficient of the first variable in the second block reflects how a general rise in single-name credit risk *interacts* with the different average correlations and correlation structures behind the two capital measures. In turn, the coefficient of the last explanatory variable captures the component of the correlation dispersion effect that is driven by a systematic relationship between individual firms' riskiness and their dependence on the common factor.

The regression results, reported in table 5, reveal that the correlation level and dispersion variables have strong explanatory power. Depending on the portfolio size, these variables explain one-third or more of the variation in calibration-driven capital discrepancies across simulated portfolios and enter the regressions with statistically significant coefficients of the expected signs. First, given that the correlation underpinning the shortcut measure stays constant across simulated portfolios, the positive impact of a higher average correlation on target capital translates into a negative impact on capital discrepancy. Second, the correlation dispersion variable enters with a positive coefficient because—given that the empirical relationship between PDs and asset return correlations is negative

[33]The correlation dispersion variable is calculated as the standard deviation of the common-factor loading coefficients, which are obtained from the single-factor approximation of the correlation matrix.

Table 5. Explaining Calibration-Driven Capital Discrepancies

	Regression 1			Regression 2		
	Large Portfolio	**Small Portfolio**	**Pooled Sample**	**Large Portfolio**	**Small Portfolio**	**Pooled Sample**
Constant	3.01	3.00	3.29	2.58	2.40	2.46
	(65.6)	(53.0)	(111.6)	(129.6)	(84.0)	(159.8)
Average Corr.	−0.24	−0.23	−0.24	−0.28	−0.26	−0.27
	(34.8)	(31.5)	(48.8)	(95.6)	(74.4)	(114.5)
Std. Dev. of Loading Coefficients	0.037	0.022	0.010	0.031	0.029	0.028
	(7.5)	(4.8)	(3.1)	(14.6)	(13.9)	(19.9)
Average PD				0.21	0.19	0.19
				(86.2)	(61.6)	(94.2)
Corr. (PD, Loading Coefficient)				−0.018	−0.019	−0.019
				(80.5)	(79.5)	(113.9)
Adjusted R^2	0.39	0.33	0.56	0.89	0.86	0.91

Note: t-statistics are shown in parentheses. The regression is based on 3,000 simulations for each portfolio size. The dependent variable equals shortcut capital (based on asset return correlation of 12 percent) minus target capital, net of the granularity and multifactor effects (see table 4). Loading coefficients are estimated under a one-common-factor approximation of the correlation structure of asset returns. The last regressor equals the Kendall rank correlation between PDs and loading coefficients.

and that the shortcut capital measure abstracts from correlation dispersion—shortcut capital overpredicts the target level by more when correlation dispersion is greater. In order to visualize this phenomenon, refer back to the dotted line in the bottom-right panel of figure 1, which represents the case of a negative correlation between PDs and common-factor loadings. In this plot, shortcut capital appears at zero correlation dispersion ($c = 0$), and a rise in correlation dispersion (i.e., a rise in c) translates into a downward movement of target capital along the dotted line.

The second part of the analysis reveals that the main driver of the correlation level and dispersion effects is the *interaction* between correlation coefficients and PDs in assessments of portfolio credit risk. In particular, adding the second block of explanatory variables to the regression raises the goodness-of-fit measures (adjusted R^2) by 50 percentage points (to 89 percent) for large portfolios and by 53 percentage points (to 86 percent) for small portfolios. In addition, the positive statistically significant coefficient of average PDs indicates that, although an increase in this variable raises both the target and shortcut capital measures, the effect is stronger under the higher (homogeneous) asset return correlation, underpinning the latter measure. Finally, the statistically significant coefficient of the correlation between PDs and common-factor loadings is with the expected negative sign. This is because target capital tends to increase in the correlation between PDs and loadings on the single factor—as illustrated by an upward movement *across* the lines in the bottom-right panel of figure 1—whereas the shortcut measure abstracts from this correlation.

Importantly, the regression results are extremely robust across portfolio sizes. The robustness can be seen in that the values of the goodness-of-fit measures, as well as the coefficient estimates and t-statistics, obtained in the context of large portfolios match almost exactly their small-portfolio counterparts. In a further test of the robustness of the regression results, we pool observations across the two portfolio sizes and observe that all estimates change only marginally, leaving the message of the regression analysis intact.[34]

[34]Background checks reveal that the residuals of the full regressions can be attributed to a large extent to interactions among PDs and asset return correlations that are nonlinear and difficult to pin down.

4.3 Estimation Errors

The above results show that capital measures based on shortcut input estimates can deviate substantially from the target level. In practice, shortcut measures are likely to be adopted by less-sophisticated users of the ASRF model who face severe constraints in terms of data and analytical capacity. By contrast, larger and more-sophisticated users are likely to construct their own estimates of asset return correlations on the basis of in-house data. This section demonstrates that, for realistic sizes of such data, small-sample estimation errors in the correlation parameters are likely to lead to large flaws in assessments of portfolio credit risk.

In order to quantify plausible estimation errors, we consider a portfolio whose "true" credit-risk parameters match those of the "typical" portfolio in our data set. For this portfolio, we impose the simplifying assumption of homogeneous PDs (1 percent), LGDs (50 percent), and pairwise asset return correlations (9.78 percent) and consider different numbers of underlying exposures (see table 6). Abstracting from issues related to granularity and multiple factors, this assumption allows us to use the ASRF model, which implies that the desired capital buffer, dubbed "benchmark," equals 3.31 percent for each portfolio size. Referring to table 4, this is recognized as the average value of the target capital buffer examined in section 4.1 above.[35]

Then, we place ourselves in the shoes of a model user who does not know the exact asset return correlations but estimates them from available data.[36] Specifically, we endow the user with 60, 120, or 300 months of asset returns data—drawn from the true underlying distribution—and calculate the sample correlation matrix. In order to quantify a plausible range of errors in the estimate of the correlation matrix, we repeat this exercise 1,000 times. As reported in panels A and B of table 6, the sample correlations contain estimation errors that are likely to be substantial even for 300 months (or twenty-five years) of data.

---

[35] Given that we abstract from model misspecification in this subsection, the benchmark capital measure is conceptually equivalent to what we earlier called target capital.

[36] In order to focus on issues in the estimation of the interdependence of credit risk across exposures, we assume that the user knows the true PD and LGD.

Table 6. Impact of Estimation Errors (in percent)

A. Sample Average of Pairwise Correlations				
	$N = 100$	$N = 200$	$N = 500$	$N = 1{,}000$
$T = 60$	9.72 [6.5, 13.3]	9.67 [6.4, 13.3]	9.64 [6.6, 13.0]	9.63 [6.9, 12.7]
$T = 120$	9.77 [7.4, 12.4]	9.77 [7.6, 12.1]	9.76 [7.6, 12.1]	9.72 [7.7, 12.0]
$T = 300$	9.75 [8.3, 11.3]	9.79 [8.3, 11.3]	9.74 [8.4, 11.2]	9.77 [8.4, 11.25]
B. Sample Standard Deviation of Loading Coefficients				
	$N = 100$	$N = 200$	$N = 500$	$N = 1{,}000$
$T = 60$	12.11 [10.3, 14.1]	11.84 [10.5, 13.2]	11.65 [10.8, 12.5]	11.58 [10.8, 12.2]
$T = 120$	8.50 [7.4, 9.8]	8.34 [7.5, 9.1]	8.17 [7.6, 8.8]	8.13 [7.7, 8.6]
$T = 300$	5.35 [4.6, 6.2]	5.24 [4.7, 5.7]	5.16 [4.8, 5.5]	5.12 [4.9, 5.4]
C. Estimated Capital, Based on One-Factor Loading Structure				
	$N = 100$	$N = 200$	$N = 500$	$N = 1{,}000$
$T = 60$	3.87 [2.8, 5.1]	3.83 [2.8, 5.0]	3.80 [2.8, 4.9]	3.79 [2.9, 4.8]
$T = 120$	3.59 [2.9, 4.4]	3.58 [2.9, 4.4]	3.56 [2.9, 4.3]	3.54 [2.9, 4.3]
$T = 300$	3.41 [3.0, 3.9]	3.42 [3.0, 3.9]	3.40 [3.0, 3.9]	3.41 [3.0, 3.9]
Benchmark	3.31	3.31	3.31	3.31

Note: Results are based on 1,000 simulations of the asset returns of N firms over T months, with the true pairwise correlation, PD, and LGD fixed at 9.78 percent, 1 percent, and 50 percent, respectively. Each cell contains the mean and the 95 percent confidence interval of the average of sample estimates of pairwise correlations (panel A), the standard deviation of the loading coefficients in a one-factor approximateion (panel B), and the implied capital measure (panel C). The "Benchmark" row in panel C refers to the level of the capital measure obtained on the basis of the true risk parameters.

Panel C of this table reveals how estimation errors in correlation coefficients translate into deviation from the desired benchmark capital buffer. First, these deviations are affected little by the number of exposures in the portfolio. Second, at standard confidence levels, the deviations decrease in the size of the available time series of asset returns but remain substantial even if this size is assumed to be unrealistically large. For example, if a portfolio comprises 1,000 exposures and a user has 120 months of data, estimated capital buffers can deviate from the benchmark level by as much as 30 percent with a 95 percent probability. For longer time series, covering 300 months, the 95 percent confidence interval does become much tighter but remains consistent with errors as high as 18 percent of the benchmark capital level. Third, estimated capital buffers exhibit a positive bias relative to the benchmark level; i.e., their average level is invariably higher than 3.31 percent. This is because the true correlation structure is assumed to be homogeneous, while small-sample errors introduce dispersion in estimated correlation coefficients. By the intuition presented in section 2.3, this dispersion raises the implied capital buffer in the presence of homogeneous PDs.

4.4 Alternative Asset Return Distributions

There is general consensus in the literature that the distribution of asset returns have tails that are fatter than the tails of the convenient normal distribution. Importantly, an erroneous normality assumption tends to bias capital buffers downward to the extent that the empirical distribution of asset returns is driven by fat tails in the distribution of the common factor.[37] Such a distribution of the common factor implies great tail dependence among asset returns, which leads to a large probability of default clustering.

In order to quantify the impact of alternative asset return distributions on capital measures, we consider a homogeneous portfolio in which all PDs equal 1 percent, all LGDs equal 50 percent, and all asset return correlations equal 9.78 percent (the same as in section 4.3). Given these risk parameters, we follow the literature on

[37] For existing theoretical and empirical analysis of the treatment of tail dependencies by credit-risk models, see, e.g., Gordy (2000), Lucas et al. (2002), and Frey and McNeil (2003).

the pricing of portfolio credit risk (see Hull and White 2004; Kalemanova, Schmid, and Werner 2007) and consider the case in which both the common and idiosyncratic factors of asset returns have the same Student-t distribution. Experimenting with different distributional specifications, we do see that fatter tails of the distribution of the common factor (i.e., fewer degrees of freedom) translate into larger deviations from a capital buffer derived under the normality assumption (table 7, left panel).

In order to examine which distributional specification is supported by the data, we rely on time series of asset returns estimated by Moody's KMV. For each of the 10,891 firms in our sample, we use the available fifty-nine months of estimated returns (from September 2001 through July 2006) to calculate the sample kurtosis, which is the standard measure of tail fatness. The mean and median of this statistic across firms equal 7.96 and 5.28, respectively. Then, on the basis of 10,000 Monte Carlo simulations, we derive the distributions of the estimators of these mean and median when (i) the data size matches the size of the Moody's KMV data on asset return estimates and (ii) both the common and idiosyncratic factors of asset returns follow the same Student-t distribution. As revealed by the confidence intervals for these estimators, the Moody's KMV data support only the "double-t" specification with 3 degrees of freedom for each factor (table 7, center and right panels).

If the asset returns are indeed driven by such distributions—which would be in line with findings in Kalemanova, Schmid, and Werner (2007)—then a normality assumption will lead to tremendous underpredictions of the desired capital buffer. Specifically, we find that a double-t specification with 3 degrees of freedom implies a capital buffer of 6.17 percent per unit of aggregate exposure. This buffer is 86 percent higher than the capital buffer calculated under a normality assumption.

Even though this result may be undermined by probable errors in the Moody's KMV estimates of asset returns, alternative double-t specifications studied in the related literature also lead to capital measures that are significantly higher than those implied by a normal distribution. Indeed, given that the available time series of Moody's KMV asset return estimates are short, plausible *systemic* errors in these estimates across firms might affect substantially the cross-sectional mean and median of the sample kurtosis. This casts

Table 7. Alternative Distributional Assumptions

Degrees of Freedom[a]	Capital Measure[b] (in percent)	Mean Kurtosis[c]		Median Kurtosis[c]	
		Mean	95% Interval	Mean	95% Interval
(3, 3)	6.17	7.83	[7.30, 8.27]	5.54	[5.11, 6.54]
(4, 4)	4.59	5.72	[5.47, 5.92]	4.43	[4.25, 4.59]
(5, 5)	4.03	4.79	[4.64, 4.93]	3.96	[3.85, 4.06]
(6, 6)	3.81	4.31	[4.19, 4.40]	3.70	[3.63, 3.78]
(7, 7)	3.67	4.02	[3.93, 4.10]	3.55	[3.49, 3.60]
(8, 8)	3.59	3.83	[3.76, 3.89]	3.44	[3.39, 3.49]
(9, 9)	3.53	3.70	[3.64, 3.75]	3.36	[3.32, 3.40]
(10, 10)	3.49	3.60	[3.55, 3.65]	3.30	[3.26, 3.34]
(∞, ∞)	3.31	3.00	[2.99, 3.01]	2.89	[2.88, 2.90]

[a]This panel shows the degrees of freedom of the Student-t distribution of asset returns' common and idiosyncratic factors, respectively. The last row refers to a Gaussian specification.
[b]Obtained by applying the general ASRF formula (equation (4)) to a homogeneous portfolio, in which PD = 1 percent and LGD = 50 percent for each exposure and asset return correlations equal 9.78 percent for each pair of exposures. The default boundary for such exposures is calculated on the basis of 10 million simulations.
[c]This panel shows the means and the 95 percent confidence intervals (based on 10,000 Monte Carlo simulations) of the estimators of the cross-sectional mean and median of the kurtosis of asset returns when (i) the marginal Student-t distributions of the common and idiosyncratic factors have the degrees of freedom specified in the row heading and asset return correlations equal 9.78 percent, and (ii) the available data comprise time series of fifty-nine returns for 10,891 firms.

doubt on the validity of the double-t specification with 3 degrees of freedom and prompts us to consider alternative specifications, with 4 degrees of freedom (which is recommended by Hull and White 2004) and 5 degrees of freedom (which is reportedly a market standard). As revealed in table 7, these alternatives imply capital measures that are, respectively, 39 and 22 percent higher than the measure incorporating normal distributions.

5. Concluding Remarks

In this paper, we have quantified the relative importance of alternative sources of error in portfolio credit-risk measures based on the popular ASRF model. Our data have revealed that violations of key modeling assumptions—namely, that granularity is perfect and that there is a single common risk factor—are likely to have a limited impact on such measures, especially for large, well-diversified portfolios. By contrast, erroneous calibration of the ASRF model—driven by flaws in popular rule-of-thumb values of asset return correlation, plausible small-sample estimation errors, or a wrong assumption regarding the distribution of asset returns—has the potential to affect substantially measures of portfolio credit risk.

Given these results and the fact that we have abstracted from several additional sources of error in portfolio credit-risk measures, the task of risk managers and supervisors appears challenging. In particular, we have assumed that PDs and LGDs are free of estimation noise. However, such noise is likely to be sizable in practice and to interact with noise in correlation estimates in generating errors in measured portfolio credit risk. In addition, there may be time variation in credit-risk parameters that relate to the likelihood and severity of default losses as well as to the correlation of the occurrence of such losses across exposures. Such time variation, which could be due either to cyclical developments or to structural changes in credit markets, would impair the useful content of the available data and, thus, would make it even more difficult to measure portfolio credit risk.

That said, the ASRF model is by no means the only way to measure portfolio credit risk (see Gordy 2000). It is, thus, important to study the degree to which specification and calibration errors affect implications of alternative models as well. Such a generalization of the analysis in this paper would be valuable to risk managers and supervisors.

Appendix 1. Gaussian Copula

The Gaussian copula is an efficient algorithm for measuring portfolio credit risk when a portfolio consists of a finite number of exposures, the correlation matrix is driven by a common-factor loading structure, and underlying distributions are normal. The efficiency of the algorithm stems from the fact that, conditional on the realization of the common factor(s), default occurrences are independent across exposures. This allows for a closed-form solution for the conditional probability of joint defaults. The corresponding unconditional probability is then derived by integrating over the probability distribution of the common factor(s). For further detail, see Gibson (2004).

Appendix 2. Monte Carlo Simulations

Monte Carlo simulations deliver the target capital level. This method can be applied to any portfolio comprising N equally weighted exposures, provided that the exposure-specific probabilities of default, PD_i; the distribution of LGD_i; and the correlation matrix of asset returns, R; are known.

Given that LGD_i is assumed to be independent of the factors underlying PD_i, and that the distribution of LGD_i is identical across exposures, the simulation of portfolio credit losses can be divided into two parts. The first part calculates the probability distribution of joint defaults. Given this distribution, the second part incorporates the LGD distribution to derive the probability distribution of portfolio losses.

Specifically, drawing on section 2.1, we estimate the probability of joint defaults as follows:

1. Using the vector $\{PD_i\}_{i=1}^{N}$ and the assumption that asset returns are distributed as standard normal variables, we obtain an $N \times 1$ vector of default thresholds.
2. We draw an $N \times 1$ 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.

3. We repeat the previous step 500,000 times to derive the probability distribution of the number of defaults, $\Pr(nd = k)$, where nd refers to the number of defaults and $k = 0, 1, \cdots, N$.

Then, we estimate the probability distribution of portfolio credit losses as follows:

1. For a given number of defaults, k, we draw LGDs for the defaulted exposures 1,000 times and calculate the conditional loss distribution, $\Pr(TL|nd = k)$.
2. We conduct the above exercise for each $k = 1, \ldots, N$, and then calculate the unconditional probability distribution of portfolio credit losses. Specifically, $\Pr(TL) = \sum_k \Pr(TL|nd = k) \cdot \Pr(nd = k)$.

Finally, we set the capital measure to equal $TL_{1-\alpha} - \sum_{i=1}^{N} E(PD_i) \cdot E(LGD_i)$.

Appendix 3. Fitting a Single-Factor Correlation Structure

A single-factor approximation of an empirical correlation matrix is obtained as follows. Denote the empirical correlation matrix by Σ and its elements σ_{ij}, for $i, j \in \{1, \cdots, N\}$. The single-factor loading structure $\rho \equiv [\rho_1, \ldots \rho_N]$ that minimizes the discrepancies between the elements of Σ and their fitted counterparts are given by

$$\min_{\rho} \sum_{i=1}^{N-1} \sum_{j>i} \left(\sigma_{ij} - \rho_i \rho_j\right)^2.$$

Andersen, Sidenius, and Basu (2003) propose an efficient algorithm to solve this minimization problem. The fitted correlation matrix $\hat{\Sigma}$ has elements $\rho_i \rho_j$.

We also construct a measure that reflects the "explanatory power" of the single-factor approximation:

$$\text{Goodness-of-Fit Measure } \equiv 1 - \frac{var(\epsilon)}{var(\sigma)},$$

where σ is a vector of all pairwise correlation coefficients σ_{ij} $(i, j = 1, \cdots, N, i < j)$ and ϵ is a vector of the errors $\sigma_{ij} - \rho_i \rho_j$ $(i, j = 1, \cdots, N, i < j)$. This measure reflects the degree to which the cross-sectional variation in pairwise correlations can be explained by common-factor loadings in a single-factor framework.

References

Andersen, L., J. Sidenius, and S. Basu. 2003. "All Your Hedges in One Basket." *Risk* 16 (November): 67–72.

Arora, N., J. Bohn, and I. Korablev. 2005. "Power and Level Validation of the EDF™ Credit Measure in the U.S. Market." Moody's KMV White Paper.

Basel Committee on Banking Supervision. 2005. "An Explanatory Note on the Basel II IRB Risk Weight Functions." BCBS Technical Note, Bank for International Settlements.

———. 2006. "Studies on Credit Risk Concentration: An Overview of the Issues and a Synopsis of the Results from the Research Task Force Project." BCBS Research Task Force Working Paper, Bank for International Settlements.

Berndt, A., R. Douglas, D. Duffie, M. Ferguson, and D. Schranz. 2005. "Measuring Default Risk Premia from Default Swap Rates and EDFs." BIS Working Paper No. 173.

Browne, C., R. Martin, and K. Thompson. 2001. "Taking to the Saddle." *Risk* 14 (June): 91–94.

Crosbie, P. 2005. "Global Correlation Factor Structure: Modeling Methodology." Moody's KMV Documents.

Crosbie, P., and J. Bohn. 2003. "Modeling Default Risk." Moody's 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* 19 (October).

Düllmann, K. 2006. "Measuring Business Sector Concentration by an Infection Model." Deutsche Bundesbank Discussion Paper Series 2, No. 2006/03.

Düllmann, K., and N. Masschelein. 2006. "Sector Concentration in Loan Portfolios and Economic Capital." Deutsche Bundesbank Discussion Paper Series 2, No. 2006/09.

———. 2007. "A Tractable Model to Measure Sector Concentration Risk in Credit Portfolios." *Journal of Financial Services Research* 32 (1): 55–79.

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* 7 (2): 85–101.

Frey, R., and A. J. McNeil. 2003. "Dependent Defaults in Models of Portfolio Credit Risk." *Journal of Risk* 6 (1): 59–92.

Garcia Cespedes, J. C., J. Antonio de Juan Herrero, A. Kreinin, and D. Rosen. 2006. "A Simple Multifactor 'Factor Adjustment' for the Treatment of Credit Capital Diversification." *Journal of Credit Risk* 2 (3): 57–85.

Gibson, M. S. 2004. "Understanding the Risk of Synthetic CDOs." Finance and Economics Discussion Series, No. 2004-36, Board of Governors of the Federal Reserve System.

Gordy, M. B. 2000. "A Comparative Anatomy of Credit Risk Models." *Journal of Banking and Finance* 24 (1–2): 119–49.

———. 2003. "A Risk-Factor Model Foundation for Ratings-Based Bank Capital Rules." *Journal of Financial Intermediation* 12 (3): 199–232.

Gordy, M., and E. Lütkebohmert. 2007. "Granularity Adjustment for Basel II." Deutsche Bundesbank Discussion Paper No. 01/2007.

Gupta, S., A. McLeod, P. Teklos, and K. Thompson. 2005. "Time for Multi-Period Capital Models." *Risk* 18 (October): 74–78.

Hanson, S. G., M. H. Pesaran, and T. Schuermann. 2007. "Firm Heterogeneity and Credit Risk Diversification." Forthcoming in *Journal of Empirical Finance.*

Heitfield, E., S. Burton, and S. Chomsisengphet. 2006. "Systematic and Idiosyncratic Risk in Syndicated Loan Portfolios." *Journal of Credit Risk* 2 (3): 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* 12 (Winter): 8–23.

Kalemanova, A., B. Schmid, and R. Werner. 2007. "The Normal Inverse Gaussian Distribution for Synthetic CDO Pricing." *Journal of Derivatives* 14 (Spring).

Kupiec, P. 2008. "A Generalized Single Common Factor Model of Portfolio Credit Risk." *Journal of Derivatives* 15 (Spring): 25–40.

Loeffler, G. 2003. "The Effects of Estimation Error on Measures of Portfolio Credit Risk." *Journal of Banking and Finance* 27 (8): 1427–53.

Longstaff, F. A., S. Mithal, and E. Neis. 2005. "Corporate Yield Spreads: Default Risk or Liquidity? New Evidence from the Credit Default Swap Market." *Journal of Finance* 60 (5): 2213–53.

Lopez, J. A. 2004. "The Empirical Relationship between Average Asset Correlation, Firm Probability of Default and Asset Size." *Journal of Financial Intermediation* 13 (2): 265–83.

Lucas, A., P. Klaassen, P. Spreij, and S. Straetmans. 2002. "Extreme Tails for Linear Portfolio Credit Risk Models." In *Risk Measurement and Systemic Risk: Proceedings of the 3rd Joint Central Bank Research Conference.* Basel, Switzerland: Bank for International Settlements.

Martin, R., and T. Wilde. 2002. "Unsystematic Credit Risk." *Risk* 15 (November): 123–28.

Merton, R. C. 1974. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates." *Journal of Finance* 29 (2): 449–70.

Morinaga, S., and Y. Shiina. 2005. "Underestimation of Sector Concentration Risk by Mis-assignment of Borrowers." Working Paper.

Ordovas, R., and K. Thompson. 2003. "Credit Ensembles." *Risk* 16 (April): 67–72.

Pykhtin, M. 2004. "Multi-factor Adjustment." *Risk* 17 (March): 85–90.

Tarashev, N., and H. Zhu. 2006. "The Pricing of Portfolio Credit Risk." Bank for International Settlements Working Paper No. 214.

Vasicek, O. 1991. "Limiting Loan Loss Probability Distribution." KMV Working Paper.

———. 2002. "Loan Portfolio Value." *Risk* 15 (December): 160–62.