Granularity Adjustment for Regulatory Capital Assessment*

Michael B. Gordy[a] and Eva Lütkebohmert[b]
[a]Federal Reserve Board
[b]University of Freiburg

The credit value-at-risk model underpinning the internal ratings-based approach of Basel II and III assumes that idiosyncratic risk has been fully diversified in the portfolio, so that economic capital depends only on systematic risk contributions. We propose a simple *granularity adjustment* (GA) for approximating the effect of undiversified idiosyncratic risk on required capital. To mitigate operational burden in implementation, we derive upper and lower bounds on the GA under incomplete information on the portfolio. We assess the magnitude and accuracy of the proposed GA on a set of bank portfolios drawn from the German credit register.

JEL Codes: G32, G28, G17.

1. Introduction

In the portfolio risk-factor frameworks that underpin both industry models of credit value-at-risk (VaR) and the internal ratings-based (IRB) risk weights of Basel II and Basel III, credit risk in a portfolio arises from two sources, systematic risk and idiosyncratic

---

*This paper originally circulated under the title "Granularity Adjustment for Basel II." Much of this work was completed while M. Gordy was a visiting scholar at the Indian School of Business and while E. Lütkebohmert was at the Deutsche Bundesbank and University of Bonn. We thank Klaus Düllmann, Chris Finger, Christian Schmieder, Dirk Tasche, and Birgit Uhlenbrock for helpful comments and input, and Jim Marrone for excellent research assistance. The opinions expressed here are our own and do not reflect the views of the Deutsche Bundesbank or of the Board of Governors of the Federal Reserve System. Author e-mails: michael.gordy@frb.gov and eva.luetkebohmert@wvf.uni-freiburg.de.

risk. Systematic risk arises because of unexpected changes in macroeconomic conditions and financial market conditions to which most borrowers are exposed. This risk cannot be eliminated through diversification across borrowers. All remaining sources of risk are idiosyncratic, i.e., particular to individual borrowers. As a portfolio becomes more fine grained, in the sense that the largest individual exposures account for a smaller share of total portfolio exposure, idiosyncratic risk is diversified away at the portfolio level.

Under the asymptotic framework that underpins the IRB approach, it is assumed that bank portfolios are *perfectly* fine grained, that is, that idiosyncratic risk has been diversified away, so that economic capital depends only on systematic risk. It is also assumed that there is only a single systematic source of risk, so that under mild regularity conditions capital charges can be calculated analytically. Relative to Monte Carlo simulation, simple closed-form capital rules are preferred in a regulatory setting for reasons of transparency, verifiability, and ease of implementation across institutions of varying capacity. Real-world portfolios are not, of course, perfectly fine grained. The asymptotic assumption might be approximately valid for some of the largest bank portfolios, but could be much less satisfactory for portfolios of smaller or more specialized institutions. When there are material name concentrations of exposure, there will be a residual of undiversified idiosyncratic risk in the portfolio. The IRB formula omits the contribution of this residual to required economic capital.

The impact of undiversified idiosyncratic risk on portfolio capital requirements can be assessed via a methodology known as *granularity adjustment*. In this paper, we propose and empirically evaluate a granularity adjustment (GA) suitable for application by banks subject to IRB capital requirements and by supervisors of such banks. Our methodology is similar in form and spirit to the GA introduced in a 2001 draft of Basel II, but exploits theoretical advances over the past decade.

In practical application, it is the data inputs (and not the formulae applied to those inputs) that can pose the most serious obstacles to cost-effective implementation. The data inputs to our GA are drawn from quantities already required for the calculation of IRB capital charges and loan-loss reserve requirements, but with one important caveat. When a bank has multiple exposures to the same underlying borrower, it is required that these multiple exposures be aggregated into a single exposure for the purpose of calculating GA inputs. For the purpose of calculating IRB capital requirements, by contrast, the identity of the borrower is immaterial, as capital charges depend only on *characteristics* of the loan and borrower (e.g., type of loan, default probability, maturity) and not on the *identity* of the borrower per se. This is a great convenience when data on different sorts of exposures are held on different computer systems, as the job of calculating capital may be delegated to those individual systems and reported back as sub-portfolio aggregates which can then be added up in a straightforward fashion to arrive at the bank-level capital and loan-loss reserve requirements. When we measure granularity, we cannot ignore borrower identity. From the perspective of single name concentration, ten loans to a single borrower together carry much more idiosyncratic risk than the same ten loans made to ten distinct borrowers of similar risk characteristics. For many institutions, the need to aggregate information across computer systems on multiple exposures to a single borrower is the most significant challenge in implementing a granularity adjustment. In defense of this aggregation requirement, we note that such aggregation would be necessary in *any* effective measure of granularity, and so is not a drawback peculiar to the GA we propose in this paper. Furthermore, one might ask how a bank can effectively manage its name concentrations without some ability to aggregate exposures across the different activities of the bank.

To reduce the burden associated with exposure aggregation, our revised GA provides for the possibility that banks be allowed to calculate the GA on the basis of the largest exposures in the portfolio, and thereby be spared the need to aggregate data on each and every borrower. To permit such an option, regulators must be able to calculate the largest possible GA that is consistent with the incomplete data provided by the bank. Our approach, therefore, is based on an upper-bound formula for the GA as a function of complete data on the m largest exposures measured in capital contributions out of a portfolio of n loans (with $m \leq n$) and summary data on the remainder of the portfolio. As m grows towards n (i.e., as the bank provides data on a larger share of its portfolio), the upper-bound formula converges to the "whole portfolio" GA. The advantage to this approach is that the bank can be permitted to choose m in accordance with its own trade-off between higher capital charges (for m small) and higher data-collection effort (for m large) and thereby reduce implementation costs.

The regulatory origins and subsequent evolution of granularity adjustment as a methodology are reviewed in section 2. We present a brief derivation of the GA in a general model setting as a first-order asymptotic expansion of value-at-risk for a large (but finite) portfolio. When applied to the widely used CreditRisk$^+$ model of portfolio credit risk, the resulting GA formula is especially tractable. Indeed, it can be expressed as a weighted sum across borrowers of capital charges and loan-loss reserve requirements. By virtue of this linearity, it is straightforward to derive effective upper and lower bounds on the GA based on partial information for the portfolio. These bounds are presented in section 3.

In section 4, we assess the magnitude and accuracy of the GA on realistic portfolios drawn from a data set of German loans. We explore comparative statics with respect to portfolio characteristics and model parameters. As the CreditRisk$^+$ model introduces a free parameter to the GA formula that is otherwise absent from IRB framework, we also use the data set to complete the calibration. Implications for practical application are drawn in the conclusion.

2. Methodology

Granularity adjustment emerged in the risk-management literature in response to policy concerns in the development of Basel II. As the IRB approach could conceivably be applied to banks of middling size or unusual portfolio concentration, the omission of capital against undiversified idiosyncratic risk could in some cases lead to material undercapitalization. To fill this gap, granularity adjustment was introduced in an early draft of Basel II, known as the Second Consultative Paper (Basel Committee on Banking Supervision 2001, hereafter CP2), as a formal component of the minimum required capital rules of the IRB approach. This GA was obtained by fitting a functional form to the gap between actual VaR and asymptotic

VaR for a sample of synthetic portfolios. Wilde (2001) suggested the true formula for the GA, which happened to coincide almost exactly with the ad hoc specification proposed in CP2. Martin and Wilde (2002) gave a more rigorous derivation of Wilde's formula based on theoretical work by Gouriéroux, Laurent, and Scaillet (2000). Ebert and Lütkebohmert (2011) extend the GA methodology to accommodate hedged positions in loan portfolios.[1]

As part of a broader effort to trim and simplify the proposed rules, the finalized Basel II agreement (Basel Committee on Banking Supervision 2006) removed the GA from the formal minimum capital requirement and included single name concentration as a matter for supervisory review.[2] Basel III (Basel Committee on Banking Supervision 2011) retains the internal ratings-based approach with minor modifications, as well as supervisory review of risk concentration. Thus, supervisory review offers a potential venue for application of our proposed methodology. Our version of the GA is a revision and extension of the original CP2 proposal. The GA of CP2 required a first-stage calculation in which the portfolio would be mapped to a homogeneous portfolio of similar characteristics. In our revised GA, the heterogeneous portfolio is used directly in the formula. The resulting algorithm is both simpler and more robust. Also, our methodology has been adapted to changes in the definition of regulatory capital between CP2 and the finalized Basel II document.

In section 2.1, we show how the GA is derived in a general model setting. Application to the IRB framework is described in section 2.2. Some alternative methodologies are considered in section 2.3.

[1]In recent years, granularity adjustment has been applied to option pricing (Gagliardini and Gouriéroux 2011); pricing and risk measurement of CDOs (Antonov, Mechkov, and Misirpashaev 2007), econometrics (Gouriéroux and Monfort 2009; Gagliardini, Gouriéroux, and Monfort 2012; Gouriéroux and Jasiak 2012); simulation methods (Gordy and Juneja 2010); modeling systemic risk contributions in banking systems (Tarashev, Borio, and Tsatsaronis 2010); and to adjusting single-factor VaR models for unmodeled sectoral concentration (Pykhtin 2004; Garcia Cespedes et al. 2006; Gürtler, Hibbeln, and Vöhringer 2010).

[2]In the parlance of Basel II, the minimum capital requirement formulae fall under Pillar 1 of the Accord, and supervisory review falls under Pillar 2.

2.1 A General Framework for Granularity Adjustment

In principle, granularity adjustment can be applied to *any* risk-factor model of portfolio credit risk, and so we begin with a very general framework. Let X denote the systematic risk factor. For consistency with the single-factor framework of Basel II and III and for ease of presentation, we assume that X is unidimensional. Let n be the number of positions in the portfolio, and assume that exposures have been aggregated so that there is a unique borrower for each position. For positions indexed by i, let A_i denote exposure at default (EAD), which we take as deterministic, and let U_i denote the loss rate (i.e., the default loss per unit EAD), which is stochastic. Let L_n be the portfolio loss rate,

$$L_n = \sum_{i=1}^{n} s_i U_i, \tag{1}$$

where s_i denotes the portfolio share of each instrument $s_i = A_i / \sum_{j=1}^{n} A_j$. We assume that conditional expected loss $\mu(x) \equiv \mathrm{E}[L_n | X = x]$ is increasing in x.

Let $\alpha_q(Y)$ denote the q^{th} percentile of the distribution of a random variable Y, so that value at risk (VaR) at the q^{th} percentile can be written as $\alpha_q(L_n)$. Under mild regularity conditions,

$$|\alpha_q(L_n) - \mathrm{E}[L_n | X = \alpha_q(X)]| \to 0 \tag{2}$$

as n grows large (see Gordy 2003, proposition 5, for a formal treatment). Loosely speaking, $\alpha_q(L_n)$ converges to the conditional expectation $\mu(\alpha_q(X)) = \mathrm{E}[L_n | X = \alpha_q(X)]$. The gap between $\alpha_q(L_n)$ and its limit $\mu(\alpha_q(X))$ is attributable to undiversified idiosyncratic risk in the portfolio. It cannot be obtained in analytical form, but can be approximated via Taylor-series expansion. We follow the derivation of Voropaev (2011).

Consider a random variable Y with continuous probability density function (pdf) $f(y)$. Let Z be another random variable with pdf conditional on $Y = y$ given by $g(z|y)$. We want to determine the quantile $\alpha_q(W)$ of the sum $W = Y + Z$. The pdf f^* of W is the convolution

$$f^*(w) = \int_{-\infty}^{\infty} f(w - z) g(z | w - z) dz. \tag{3}$$

Applying a Taylor expansion of $(w-z)$ around w to the integrand yields

$$f(w-z)g(z|w-z) = \sum_{i=0}^{\infty} \frac{1}{i!} \left[\frac{d^i}{dw^i} f(w)g(z|w) \right] ((w-z)-w)^i$$
$$= f(w)g(z|w) + \sum_{i=1}^{\infty} \frac{(-1)^i}{i!} \left[\frac{d^i}{dw^i} f(w)g(z|w) \right] z^i.$$

Substituting into (3) and applying Fubini's theorem, we have

$$f^*(w) = \int_{-\infty}^{\infty} f(w)g(z|w)dz + \sum_{i=1}^{\infty} \frac{(-1)^i}{i!} \frac{d^i}{dw^i} \left[\int_{-\infty}^{\infty} f(w)g(z|w)z^i dz \right]. \tag{4}$$

Let $m_i(w)$ denote the i^{th} moment of Z conditional on $Y=w$, i.e.,

$$m_i(w) = \mathrm{E}[Z^i|Y=w] = \int_{-\infty}^{\infty} z^i \, g(z|w) \, dz$$

and substitute

$$\int_{-\infty}^{\infty} f(w)g(z|w)z^i dz = f(w) \int_{-\infty}^{\infty} g(z|w)z^i dz = f(w)m_i(w)$$

into (4) to arrive at

$$f^*(w) = f(w) + \sum_{i=1}^{\infty} \frac{(-1)^i}{i!} \frac{d^i}{dw^i} [f(w)m_i(w)]. \tag{5}$$

By construction, the quantiles $\alpha_q(Y+Z)$ and $\alpha_q(Y)$ satisfy

$$\int_{-\infty}^{\alpha_q(Y)} f(w)dw = q = \int_{-\infty}^{\alpha_q(Y+Z)} f^*(w)dw. \tag{6}$$

Inserting equation (5) into (6) gives

$$\int_{\alpha_q(Y)}^{\alpha_q(Y+Z)} f(w)dw = \int_{-\infty}^{\alpha_q(Y+Z)} f(w)dw - \int_{-\infty}^{\alpha_q(Y)} f(w)dw$$
$$= \int_{-\infty}^{\alpha_q(Y+Z)} f^*(w)dw$$
$$- \int_{-\infty}^{\alpha_q(Y+Z)} \sum_{i=1}^{\infty} \frac{(-1)^i}{i!} \frac{d^i}{dw^i}[f(w)m_i(w)]dw - \int_{-\infty}^{\alpha_q(Y)} f(w)dw$$
$$= \sum_{i=1}^{\infty} \frac{(-1)^{i-1}}{i!} \frac{d^{i-1}}{dw^{i-1}}[f(w)m_i(w)]\Big|_{w=\alpha_q(Y+Z)}.$$

Expanding both sides in powers of $(\alpha_q(Y+Z) - \alpha_q(Y))$ around $\alpha_q(Y)$ and assuming for the moment that $m_1(w) = 0$, we obtain

$$f(\alpha_q(Y)) \cdot (\alpha_q(Y+Z) - \alpha_q(Y)) \approx -\frac{1}{2}\frac{d}{dw}[f(w)m_2(w)]\Big|_{w=\alpha_q(Y)} \tag{7}$$

and thus

$$\alpha_q(Y+Z) - \alpha_q(Y) \approx -\frac{1}{2f(\alpha_q(Y))} \cdot \frac{d}{dw}[f(w)m_2(w)]\Big|_{w=\alpha_q(Y)}. \tag{8}$$

For Y, substitute the asymptotic portfolio loss $\mu(X)$, and for Z substitute the residual $L_n - \mu(X)$. Since $\mu(x)$ is assumed to be strictly increasing, conditioning on $Y = \mu(x)$ is equivalent to conditioning on $X = x$. For any $w = \mu(x)$, we have

$$m_1(w) = \mathrm{E}[Z|X=x] = \mathrm{E}[L_n|X=x] - \mathrm{E}[\mathrm{E}[L_n|X=x]|X=x] = 0,$$

as required for equation (7), and

$$m_2(w) = \mathrm{E}[Z^2|X=x] = \mathrm{V}[L_n|X=x] \equiv \sigma^2(x).$$

That is, $\sigma^2(x)$ is the conditional variance of portfolio loss. We substitute into equation (8) to obtain the granularity adjustment as

$$\begin{aligned}\alpha_q(L_n) - \mu(\alpha_q(X)) \approx {} & -\frac{1}{2f(\mu(\alpha_q(X)))} \\ & \times \frac{d}{d\mu(x)}[f(\mu(x))m_2(\mu(x))]\bigg|_{\mu(x)=\alpha_q(\mu(X))} \\ = {} & -\frac{1}{2h(\alpha_q(X))}\frac{d}{dx}\left[\frac{h(x)\sigma^2(x)}{\mu'(x)}\right]\Bigg|_{x=\alpha_q(X)} \equiv \text{GA},\end{aligned} \tag{9}$$

where h is the pdf of X and we have used $f(w)dw = \mu'(x)h(x)dx$.

We have so far imposed no specific accounting definition of "loss," so we can choose to measure the $\{U_i\}$ on a mark-to-market basis or an actuarial basis, and either inclusive or exclusive of expected loss. The latter point is important in light of the change in the definition of capital between CP2 and the final Basel II document. CP2 tied capital requirements to a quantile of the portfolio loss distribution, i.e., on VaR. In the final version of Basel II and in Basel III, the VaR requirement is decomposed into expected loss (EL) and unexpected loss (UL) components. Expected loss provides the basis for the loan-loss reserve requirement. Unexpected loss is simply defined as the difference between VaR and EL, and provides the basis for the minimum capital requirement. Let UL_n be the unexpected loss $\alpha_q(L_n) - \text{E}[L_n]$, and observe that

$$|\text{UL}_n - (\mu(\alpha_q(X)) - \text{E}[L_n])| = |\alpha_q(L_n) - \mu(\alpha_q(X))| \to 0$$

by equation (2). Therefore, $\mu(\alpha_q(X)) - \text{E}[L_n]$ is the asymptotic capital charge under a UL regime. The granularity adjustment for UL is the same as the granularity adjustment for VaR, so the change in the definition of regulatory capital leaves the GA unchanged. However, as we will see below, the *inputs* to the GA formula do depend on the distinction between EL and UL. In revising the GA, we have adapted the formula for consistency with the finalized Basel II and III rules.

Although value-at-risk is ubiquitous in industry practice and is the risk measure embedded in Basel II and III, it is well understood that it has theoretical and practical shortcomings (see, e.g., McNeil, Frey, and Embrechts 2005, §6.1). A popular alternative to VaR is *expected shortfall*. In the appendix, we present the granularity adjustment for this alternative, as well as the upper and lower bounds corresponding to those of section 3.

2.2 A Granularity Adjustment for Basel II and III

In the GA formula, the expressions for $\mu(x)$, $\sigma^2(x)$, and $h(x)$ are model dependent. For application of the GA in a regulatory setting, it would be desirable to base the GA on the same model as that which underpins the IRB capital formula in order to avoid internal inconsistencies. Unfortunately, this would be difficult to implement. The IRB formula is derived within a single-factor mark-to-market Vasicek model closest in spirit to KMV Portfolio Manager. Gordy and Marrone (2012) show how to construct $\mu(x)$ and $\sigma^2(x)$ for models in this class, but the resulting expressions lack the parsimony and transparency that are desirable for supervisory application.[3]

As fidelity to the IRB model cannot readily be achieved by direct means, we adopt an indirect strategy. We base the GA on a model chosen for the tractability of the expressions appearing in the general formula (9), and then reparameterize the inputs of the formula in a way that restores consistency with the IRB model in loan-level reserve requirements and capital charges. Our chosen model is an extended version of the single-factor CreditRisk$^+$ model that allows for idiosyncratic recovery risk. CreditRisk$^+$ is a widely used industry model for portfolio credit risk that was introduced by Credit Suisse Financial Products (1997). A variety of extensions can be found in Gundlach and Lehrbass (2004).

As CreditRisk$^+$ is an actuarial model of loss, we define the loss rate as $U_i = \mathrm{LGD}_i \cdot D_i$, where D_i is a default indicator equal to 1 if the borrower defaults, 0 otherwise, and where LGD_i is the loss rate in the event of default on loan i. The systematic factor X generates

[3]A further barrier to using this model is that the IRB formula is not fit to the model directly, but rather is linearized with respect to maturity. The “true” term structure of capital charges in mark-to-market models tends to be strongly concave, so this linearization is not at all a minor adjustment. It is not clear how one would alter $\mu(x)$ and $\sigma^2(x)$ to make the GA consistent with the linearized IRB formula.

correlation across borrower defaults by shifting the default probabilities. Conditional on $X = x$, defaults are independent across borrowers and occur with probability

$$\pi_i(x) = \mathrm{PD}_i \cdot (1 + \omega_i(x - 1))$$

and PD_i is the unconditional probability of default for borrower i. The factor loading ω_i controls the sensitivity of borrower i to the systematic risk factor. We assume that X is gamma-distributed with mean 1 and variance $1/\xi$, where $\xi > 0$ is the *precision parameter*. Hence, the $\{D_i\}$ are conditionally independent Bernoulli random variables with $\Pr(D_i = 1|X = x) = \pi_i(x)$, whereas unconditionally the $\{D_i\}$ are dependent Bernoulli random variables with $\Pr(D_i = 1) = \mathrm{PD}_i$.[4] Finally, to obtain an analytical solution to the model, in CreditRisk$^+$ one approximates the distribution of the default indicator variable as a Poisson distribution.[5]

In the standard version of CreditRisk$^+$, the recovery rate is assumed to be known with certainty. Our extended model allows LGD_i to be a random loss given default with expected value $\mathrm{E}[\mathrm{LGD}_i]$ and variance $\mathrm{V}[\mathrm{LGD}_i]$. The LGD uncertainty is assumed to be entirely idiosyncratic and therefore independent of X. This independence assumption may be somewhat restrictive in view of the evidence of Altman et al. (2005), Acharya, Bharath, and Srinivasan (2007), and others, but is retained here for consistency with model assumptions underpinning the IRB formulae as well as for tractability.

---

[4]Because X can take any positive real value, the conditional default probability can be greater than 1. For realistic PDs and empirically plausible values of ξ, the likelihood that $\pi(X) > 1$ is very small. Even in the situation of a large PD of 10 percent, a high factor loading of $\omega = 1$, and a low precision parameter value of $\xi = 0.03$, we find $\Pr(\pi(X) > 1) < 0.03$. Capping $\pi(X)$ at 1 (which sacrifices the analytical solution to the model) has a negligible effect on the loss distribution (Gordy 2000, table 7).

[5]Approximation of a Bernoulli default event as a Poisson random variable is crucial to maintaining analytical tractability when we aggregate to the portfolio loss distribution. This so-called Poisson approximation has long been used in the actuarial literature (see Panjer and Willmot 1992, §6.5). The associated approximation error is proportional to PD^2, so it is small when most borrowers have low PD. Error bounds in the context of CreditRisk$^+$ are provided by Merino and Nyfeler (2004). In a numerical example on a representative large bank portfolio, Gordy (2000, table 7) shows that the approximation error in quantiles of the loss distribution is negligible.

We next obtain the $\mu(x)$ and $\sigma^2(x)$ functions for this model. Let us define at the instrument level the functions $\mu_i(x) = \mathrm{E}[U_i|X = x]$ and $\sigma_i^2(x) = \mathrm{V}[U_i|X = x]$. By the conditional independence assumption, we have

$$\mu(x) = \mathrm{E}[L_n|X = x] = \sum_{i=1}^{n} s_i \mu_i(x)$$
$$\sigma^2(x) = \mathrm{V}[L_n|X = x] = \sum_{i=1}^{n} s_i^2 \sigma_i^2(x).$$

In an actuarial setting, the borrower's conditional expected loss function is $\mu_i(x) = \mathrm{E}[\mathrm{LGD}_i] \cdot \pi_i(x)$, and the conditional variance of loss is

$$\begin{aligned}\sigma_i^2(x) &= \mathrm{E}[\mathrm{LGD}_i^2 \cdot D_i^2|X = x] - \mathrm{E}[\mathrm{LGD}_i]^2 \cdot \pi_i(x)^2 \\ &= \mathrm{E}[\mathrm{LGD}_i^2] \cdot \mathrm{E}[D_i^2|X = x] - \mu_i(x)^2. \end{aligned} \tag{10}$$

As D_i given X is assumed in CreditRisk$^+$ to be Poisson distributed, we have $\mathrm{E}[D_i|X = x] = \mathrm{V}[D_i|X = x] = \pi_i(x)$, which implies

$$\mathrm{E}[D_i^2|X = x] = \pi_i(x) + \pi_i(x)^2.$$

For the term $\mathrm{E}[\mathrm{LGD}_i^2]$ in the conditional variance, we can substitute

$$\mathrm{E}[\mathrm{LGD}_i^2] = \mathrm{V}[\mathrm{LGD}_i] + \mathrm{E}[\mathrm{LGD}_i]^2.$$

This leads us to

$$\begin{aligned}\sigma_i^2(x) &= \left(\mathrm{V}[\mathrm{LGD}_i] + \mathrm{E}[\mathrm{LGD}_i]^2\right) \cdot \left(\pi_i(x) + \pi_i(x)^2\right) - \mu_i(x)^2 \\ &= \mathcal{C}_i \mu_i(x) + \mu_i(x)^2 \cdot \frac{\mathrm{V}[\mathrm{LGD}_i]}{\mathrm{E}[\mathrm{LGD}_i]^2},\end{aligned}$$

where $\mathcal{C}_i$ is defined as

$$\mathcal{C}_i \equiv \frac{\mathrm{V}[\mathrm{LGD}_i] + \mathrm{E}[\mathrm{LGD}_i]^2}{\mathrm{E}[\mathrm{LGD}_i]}. \tag{11}$$

We substitute the expressions for $\mu(x)$ and $\sigma^2(x)$ into equation (9) and then evaluate the derivative in that equation at $x = \alpha_q(X)$ to obtain

$$\begin{aligned}
\mathrm{GA} = &-\frac{1}{2\sum_{i=1}^n s_i\mu_i'(\alpha_q(X))} \cdot \frac{h'(\alpha_q(X))}{h(\alpha_q(X))} \cdot \sum_{i=1}^n s_i^2 \\
&\times \left(\mathcal{C}_i\mu_i(\alpha_q(X)) + \mu_i^2(\alpha_q(X))\frac{\mathrm{V}[\mathrm{LGD}_i]}{\mathrm{E}[\mathrm{LGD}_i]^2}\right) \\
&- \frac{1}{2\sum_{i=1}^n s_i\mu_i'(\alpha_q(X))} \cdot \sum_{i=1}^n s_i^2\mu_i'(\alpha_q(X)) \\
&\times \left(\mathcal{C}_i + 2\mu_i(\alpha_q(X))\frac{\mathrm{V}[\mathrm{LGD}_i]}{\mathrm{E}[\mathrm{LGD}_i]^2}\right). \qquad (12)
\end{aligned}$$

Through the $\mu_i(x)$ function, this formula depends on the instrument-level parameters PD_i and ω_i, as well as on $\mathrm{E}[\mathrm{LGD}_i]$ and $\mathrm{V}[\mathrm{LGD}_i]$.

We now reparameterize the inputs. Let $\mathcal{R}_i$ be the loan-loss reserve requirement as a share of EAD for instrument i. In the actuarial setting of CreditRisk$^+$, this is simply

$$\mathcal{R}_i = \mathrm{E}[U_i] = \mathrm{E}[\mathrm{LGD}_i] \cdot \mathrm{PD}_i.$$

Let $\mathcal{K}_i$ be the asymptotic UL capital requirement as a share of EAD. In CreditRisk$^+$, this is

$$\mathcal{K}_i = \mathrm{E}[U_i|X = \alpha_q(X)] - \mathrm{E}[U_i] = \mathrm{E}[\mathrm{LGD}_i] \cdot \mathrm{PD}_i \cdot \omega_i \cdot (\alpha_q(X) - 1). \qquad (13)$$

At the portfolio level, we define $\mathcal{R}^*$ and $\mathcal{K}^*$ analogously as the required loan-loss reserve and required capital per unit exposure for the portfolio as a whole, i.e.,

$$\mathcal{R}^* = \sum_{i=1}^n s_i\mathcal{R}_i \quad \text{and} \quad \mathcal{K}^* = \sum_{i=1}^n s_i\mathcal{K}_i.$$

When we substitute $\mathcal{R}_i$ and $\mathcal{K}_i$ into the CreditRisk$^+$ GA, we find that the PD_i and ω_i inputs can be eliminated. We arrive at the following formula:

$$\text{GA} = \frac{1}{2\mathcal{K}^*}\sum_{i=1}^{n} s_i^2 \left[\left(\delta\mathcal{C}_i(\mathcal{K}_i+\mathcal{R}_i) + \delta(\mathcal{K}_i+\mathcal{R}_i)^2 \cdot \frac{\text{V}[\text{LGD}_i]}{\text{E}[\text{LGD}_i]^2}\right)\right.$$
$$\left. -\mathcal{K}_i\left(\mathcal{C}_i + 2(\mathcal{K}_i+\mathcal{R}_i)\cdot\frac{\text{V}[\text{LGD}_i]}{\text{E}[\text{LGD}_i]^2}\right)\right], \tag{14}$$

where

$$\delta \equiv -\frac{h'(x)}{h(x)}(\alpha_q(X)-1) = (\alpha_q(X)-1)\cdot\left(\xi + \frac{1-\xi}{\alpha_q(X)}\right). \tag{15}$$

Observe that the expression for δ depends only on model parameters, not data inputs, so δ is a regulatory parameter. It is through δ that the precision parameter ξ influences the GA. The GA proposal of CP2 assumes $\xi = 0.25$. With the target solvency probability set to $q = 0.999$, this value of ξ implies $\delta = 4.83$. We return to the calibration of ξ in section 4.3.

The variance of LGD neither is an input to the IRB formula, nor is it restricted in any way within the IRB model. Banks could, in principle, be permitted or required to supply this parameter for each loan. Given the limited data currently available on recoveries, it seems preferable to impose a regulatory assumption on V[LGD] in order to avoid the burden of a new data requirement. We impose the relationship as found in the CP2 version of the GA:

$$\text{V}[\text{LGD}_i] = \gamma\text{E}[\text{LGD}_i](1-\text{E}[\text{LGD}_i]), \tag{16}$$

where the regulatory parameter γ is between 0 and 1. When this specification is used in industry models such as CreditMetrics and KMV Portfolio Manager, a typical setting is $\gamma = 0.25$.

The GA formula can be simplified somewhat. The quantities $\mathcal{R}_i$ and $\mathcal{K}_i$ are typically small, and so terms that are products of these quantities can be expected to contribute little to the GA. If these second-order terms are dropped, we arrive at the simplified formula:

$$\widetilde{\text{GA}} = \frac{1}{2\mathcal{K}^*}\sum_{i=1}^{n} s_i^2\mathcal{C}_i\left(\delta(\mathcal{K}_i+\mathcal{R}_i)-\mathcal{K}_i\right). \tag{17}$$

Here and henceforth, we use the tilde to indicate this simplified GA formula, and refer to the GA of equation (14) as the "exact" GA. The accuracy of this approximation to the exact GA is evaluated in section 4.2. To summarize, we outline the implementation of $\widetilde{\mathrm{GA}}$ in the following algorithm.

Implementation of Simplified GA:

Bank Inputs: default probability PD_i, expected loss given default $\mathrm{E}[\mathrm{LGD}_i]$, exposure at default A_i and maturity M_i for each borrower $i = 1, \ldots, n$

Regulatory Parameters: quantile level q, transformed precision parameter δ, and recovery risk parameter $\gamma \in (0, 1)$.

for $i = 1, \ldots, n$ **do**

(i) Set exposure shares $s_i = A_i / \sum_{j=1}^n A_j$.

(ii) Set $\mathrm{V}[\mathrm{LGD}_i] = \gamma \mathrm{E}[\mathrm{LGD}_i](1 - \mathrm{E}[\mathrm{LGD}_i])$ and

$$\mathcal{C}_i \equiv \frac{\mathrm{V}[\mathrm{LGD}_i] + \mathrm{E}[\mathrm{LGD}_i]^2}{\mathrm{E}[\mathrm{LGD}_i]}.$$

(iii) Compute loan-loss reserve requirement $\mathcal{R}_i = \mathrm{E}[\mathrm{LGD}_i] \cdot \mathrm{PD}_i$.

(iv) Compute UL capital requirement $\mathcal{K}_i$ as a function of PD_i, $\mathrm{E}[\mathrm{LGD}_i]$, and M_i by IRB formula (Basel Committee on Banking Supervision 2011, ¶102).

Compute portfolio capital requirement $\mathcal{K}^* = \sum_{i=1}^n s_i \mathcal{K}_i$.

Compute $\widetilde{\mathrm{GA}} = \frac{1}{2\mathcal{K}^*} \sum_{i=1}^n s_i^2 \mathcal{C}_i \left(\delta(\mathcal{K}_i + \mathcal{R}_i) - \mathcal{K}_i\right)$.

2.3 Alternative Approaches

Before proceeding, we pause to mention some alternative methodologies. Perhaps the very simplest approach would be based on a Herfindahl-Hirschman Index (HHI), which is defined as the sum of the squares of the portfolio shares of the individual exposures. The HHI is a continuous measure with zero corresponding to the asymptotic case (each borrower has an infinitesimal share) and unity corresponding to monopoly (there is only one borrower). Holding all else equal, the closer the HHI of a portfolio is to 1, the more concentrated the portfolio is, so the higher the appropriate granularity add-on charge. However, in the absence of a model, it is difficult to determine what the appropriate add-on for a given HHI should be. Furthermore, the effect of granularity on economic capital is sensitive to the credit quality of the portfolio, and it may be difficult to accommodate this dependence with parsimony and accuracy in ad hoc methods of granularity adjustment based on exposure HHI. Finally, an HHI-based approach does not avoid in any way the operational burden associated with aggregation of multiple exposures to a single exposure per borrower.

Another approach, due to Vasicek (2002), lies somewhere between ad hoc and model based. In this method, one augments the systematic risk contribution (by increasing the factor loading) in order to compensate for ignoring the idiosyncratic risk. The trouble is that systematic and idiosyncratic risk have very different impacts on the shape of the loss distribution. This method is known to perform poorly for realistic portfolios.

Much closer to our proposal in spirit and methodology is the approach of Emmer and Tasche (2005), who offer a granularity adjustment based on a one-factor actuarial version of the CreditMetrics model. This has the advantage of relative proximity to the model underpinning the IRB formula, but we believe this proximity to be more in appearance than in substance because the IRB formula is based on a mark-to-market generalization of this model. Capital charges are invariant to loan maturity in the actuarial model, which is inconsistent with the first-order effect of loan maturity on IRB capital charges. Gordy and Marrone (2012) demonstrate that the GA for the mark-to-market CreditMetrics model does indeed depend on maturity, so simplification of the model to its actuarial counterpart would entail a substantive loss of fidelity. A virtue of the methodology in section 2.2 is that the effect of maturity on the GA can be captured indirectly through reparameterization, because the IRB capital charges $\mathcal{K}_i$ are dependent on loan maturity. The Emmer and Tasche (2005) formula is not amenable to such reparameterization. Furthermore, even with simplification to the actuarial model, the resulting expression for the GA (as given in Emmer and Tasche 2005, proposition 2.2) remains more complex than desirable for supervisory application.

An alternative that has not been much studied is the saddlepoint method of Martin and Wilde (2002). Results in that paper suggest that it would be quite similar to the GA in performance and pose a similar trade-off between fidelity to the IRB model and analytical tractability. Lütkebohmert (2012) finds that the saddlepoint method is computationally unwieldy for the actuarial CreditMetrics model studied by Emmer and Tasche (2005) and, even in case of the CreditRisk$^+$ model, cannot accommodate hedged positions.

Finally, it is possible to estimate a granularity adjustment by Monte Carlo simulation. The advantage to simulation is flexibility in model specification, as we can estimate the GA for the IRB model or, indeed, for more general multi-factor models. In each trial of the simulation ($i = 1, \ldots, I$), we would first simulate the vector X_i of systematic risk factors. Next, we would simulate the portfolio loss L_i conditional on X_i and would calculate (possibly by a nested simulation) the conditional portfolio expected loss $\bar{L}_i$. Portfolio VaR is estimated as the q^{th} percentile of the sample $\{L_i\}$, asymptotic VaR as the q^{th} percentile of the sample $\{\bar{L}_i\}$, and the granularity adjustment as the difference between the two estimated quantiles. Because the GA is small relative to the VaR and asymptotic VaR, many trials are needed to achieve reasonable accuracy. As noted in the Introduction, in regulatory settings there is a preference for simple closed-form formulae.

3. Bounds Based on Incomplete Data

As discussed in the Introduction, aggregation of multiple exposures into a single exposure per borrower is very likely to be the only substantive challenge in implementing the granularity adjustment. Aggregation of exposures is required for day-to-day risk-management purposes, so it might seem strange that it could be a burden in practice. Nonetheless, the Institute of International Finance (2009, p. 27) post-crisis white paper on financial services reform acknowledges that aggregation must be done manually and with difficulty at many large institutions.[6] As explained in that report, the primary challenge is that many bank information systems are "siloed" by business unit. These silos are often a legacy of

[6]The Counterparty Risk Management Policy Group III (2008, §IV.C.3) white paper acknowledges this ongoing deficiency as well.

bank growth through acquisition, and reflect the difficulty of integrating disparate data systems. In some cases, silos may reflect a management culture favoring maximum operating independence at the business-unit level. A secondary challenge is that a given borrower could appear under a variety of names in different databases. Even for corporate borrowers, there is no national database of legal entities in many markets (including the United States), so banks often rely on private vendor databases to track borrowers by unique identifiers and to track corporate hierarchies among related sets of borrowers. If different bank operating units contract with different vendors, then the bank must construct and maintain a concordance table at significant expense.

To reduce this burden on the banks, we propose that banks be permitted to calculate the GA based on a *subset* consisting of the largest exposures. An upper bound can be calculated for the influence of exposures that are left out of the computation. This approach is conservative from a regulatory point of view because the upper bound is always at least as large as the "true" GA. The bank can therefore be given the flexibility to find the best trade-off between the cost of data collection and the cost of the additional capital associated with the upper bound. In order to convey most clearly the intuition behind our approach, we first present in section 3.1 the upper bound in the special case of a portfolio that is homogeneous in PD and E[LGD]. We then present in section 3.2 the upper bound for the more realistic case of a heterogeneous portfolio.

In section 3.3, we derive a lower bound on the GA, which is useful when the GA can be computed for the wholesale banking book but not for the remaining portfolio of retail and SME (small and medium enterprise) loans. In this case, the GA for the entire portfolio is bounded below by a multiple of the GA for the wholesale portfolio. As retail and SME portfolios are typically fine grained, the lower bound can be tight in practice.

3.1 Upper Bound in the Homogeneous Case

The simplest upper bound is for the case in which exposures are homogeneous in PD and E[LGD] but heterogeneous in exposure size. Assume that the bank has determined the m largest aggregate exposures in the portfolio of n borrowers ($m \leq n$) and that we have sorted these aggregated EAD values as $A_1 \geq A_2 \geq \ldots \geq A_m$. The shares

$s_1 \geq s_2 \geq \ldots \geq s_m$ are, as in section 2, calculated with respect to the *total* portfolio EAD in the denominator.

When PD and E[LGD] are homogeneous, we have $\mathcal{K}_i = \mathcal{K}^* = \mathcal{K}$ and $\mathcal{R}_i = \mathcal{R}$ for all i, and similarly $\mathcal{C}_i = \mathcal{C}$ is also independent of i. Hence the simplified GA reads

$$\widetilde{\mathrm{GA}} = \frac{1}{2\mathcal{K}}\mathcal{C}(\delta(\mathcal{K}+\mathcal{R})-\mathcal{K}) \cdot HHI,$$

where HHI is the Herfindahl-Hirschman Index, $HHI = \sum_{i=1}^{n} s_i^2$. Using only the first $m \leq n$ exposures, and defining S_m as the cumulative share of these exposures, $S_m = \sum_{i=1}^{m} s_i$, we know that HHI is bounded by

$$\begin{aligned} HHI &= \sum_{i=1}^{m} s_i^2 + \sum_{i=m+1}^{n} s_i^2 \leq \sum_{i=1}^{m} s_i^2 + s_m \cdot \sum_{i=m+1}^{n} s_i \\ &= \sum_{i=1}^{m} s_i^2 + s_m \cdot (1 - S_m). \end{aligned}$$

This leads to the following upper bound for the simplified granularity adjustment:

$$\widetilde{\mathrm{GA}} \leq \frac{1}{2\mathcal{K}}\mathcal{C}(\delta(\mathcal{K}+\mathcal{R})-\mathcal{K}) \cdot \left(\sum_{i=1}^{m} s_i^2 + s_m \cdot (1 - S_m) \right). \qquad (18)$$

3.2 Upper Bound in the Heterogeneous Case

In the general case of a heterogeneous portfolio, the upper bound becomes more complicated because the meaning of "largest exposures" is no longer unambiguous. Do we mean largest by EAD, by capital contribution, or by some other measure? It turns out that we require information on both the distribution of aggregated positions by EAD and by capital contribution. Specifically, we assume the following:

- The bank has identified the m borrowers to whom it has the largest aggregated exposures measured in capital contribution, i.e., $s_i \cdot \mathcal{K}_i$. Denote this set of borrowers as Ω. For each borrower $i \in \Omega$, the bank knows $(s_i, \mathcal{K}_i, \mathcal{R}_i, \mathcal{C}_i)$.
- For the $n - m$ exposures that are unreported (that is, exposures for which the borrower is not in Ω), the bank determines

an upper bound $\bar{s}$ on exposure share such that $s_i \leq \bar{s}$ for all $i \notin \Omega$.
- The bank knows $\mathcal{K}^*$ and $\mathcal{R}^*$ for the portfolio as a whole.

The first assumption is straightforward and unavoidable, as this is where the need arises to aggregate multiple exposures for a subset of borrowers in the portfolio. Internal risk-management reporting typically includes a list of the "tallest trees" in capital usage by customer, and therefore it is reasonable to assume that aggregated capital contribution data for the largest customers are internally available.[7] If such data cannot be made available, we might question whether the bank is making any substantive business use of its internal economic capital models.

The second assumption is necessary in order to obtain a bound on unreported exposure shares. A bound may often be implicit in a bank's risk-management practices. For example, internal risk-management systems may report on the borrowers to which the bank has the greatest exposure in EAD, in which case $\bar{s}$ can be set to the smallest exposure share on the report.

The third assumption hardly needs justification, as these portfolio-level quantities are calculated in the course of determining IRB capital requirements. In particular, $\mathcal{K}^*$ and $\mathcal{R}^*$ can be obtained in the usual manner (i.e., as share-weighted averages of exposure-level capital charges and loan-loss reserve requirements) without aggregation of exposures by borrower.

For notational convenience, define

$$Q_i \equiv \delta(\mathcal{K}_i + \mathcal{R}_i) - \mathcal{K}_i$$

and decompose the GA as

$$\begin{aligned}\widetilde{\mathrm{GA}} &= \frac{1}{2\mathcal{K}^*}\sum_{i=1}^{n} s_i^2 \mathcal{C}_i \left(\delta(\mathcal{K}_i + \mathcal{R}_i) - \mathcal{K}_i\right) \\ &= \frac{1}{2\mathcal{K}^*}\left(\sum_{i\in\Omega} s_i^2 Q_i \mathcal{C}_i + \sum_{i\notin\Omega} s_i^2 Q_i \mathcal{C}_i\right). \qquad (19)\end{aligned}$$

[7]For the sake of concrete example, we have developed a procedure by which a central manager in a siloed information system could obtain sufficient information to satisfy our assumptions. Details are available from the authors upon request.

The summation over $i \in \Omega$ is known by assumption 1. By assumption 2, we know that $\bar{s} \geq s_i$ for $i \notin \Omega$. Our assumption on V[LGD] in equation (16) is sufficient to guarantee that $\mathcal{C}_i \leq 1$. Therefore,

$$\sum_{i \notin \Omega} s_i^2 Q_i \mathcal{C}_i \leq \bar{s} \sum_{i \notin \Omega} s_i Q_i = \bar{s} \left(\delta \sum_{i \notin \Omega} s_i (\mathcal{K}_i + \mathcal{R}_i) - \sum_{i \notin \Omega} s_i \mathcal{K}_i \right).$$

Extending the notation of section 3.1, we define $S_\Omega = \sum_{i \in \Omega} s_i$ as the share of total portfolio exposure that is in the observed set Ω. We generalize the notation $\mathcal{K}^*$ and $\mathcal{R}^*$ so that

$$\mathcal{K}^*_\Omega = \frac{1}{S_\Omega} \sum_{i \in \Omega} s_i \mathcal{K}_i \quad \text{and} \quad \mathcal{R}^*_\Omega = \frac{1}{S_\Omega} \sum_{i \in \Omega} s_i \mathcal{R}_i$$

are (respectively) the share-weighted average capital requirement and loan-loss reserve requirement for sub-portfolio Ω. Assumption 1 implies that $\mathcal{K}^*_\Omega$ and $\mathcal{R}^*_\Omega$ are known to the bank, which in turn implies that

$$\sum_{i \notin \Omega} s_i \mathcal{K}_i = \mathcal{K}^* - S_\Omega \mathcal{K}^*_\Omega$$
$$\sum_{i \notin \Omega} s_i \mathcal{R}_i = \mathcal{R}^* - S_\Omega \mathcal{R}^*_\Omega$$

are also known to the bank. Thus, we arrive at

$$\sum_{i \notin \Omega} s_i^2 Q_i \mathcal{C}_i \leq \bar{s}((\delta - 1)(\mathcal{K}^* - S_\Omega \mathcal{K}^*_\Omega) + \delta(\mathcal{R}^* - S_\Omega \mathcal{R}^*_\Omega)). \tag{20}$$

Finally, we substitute into equation (19) to obtain

$$\widetilde{\text{GA}} \leq \frac{1}{2\mathcal{K}^*} \left(\sum_{i \in \Omega} s_i^2 Q_i \mathcal{C}_i + \bar{s}\left((\delta - 1)(\mathcal{K}^* - S_\Omega \mathcal{K}^*_\Omega) + \delta(\mathcal{R}^* - S_\Omega \mathcal{R}^*_\Omega)\right) \right). \tag{21}$$

3.3 Lower Bound in the Heterogeneous Case

A lower bound on the GA can be obtained as a special case of a more general problem. When a portfolio is decomposed into two sub-portfolios, what is the relationship between the whole portfolio

GA and the stand-alone GAs for the sub-portfolios? In the spirit of our earlier notation, we partition the portfolio into Ω and its complement Ω^c. We write S_Ω for the share of exposure in Ω to total portfolio exposure and write $\mathcal{K}^*_\Omega$ for the share-weighted average capital requirement within Ω. Next, define $\widetilde{\mathrm{GA}}_\Omega$ as the simplified GA for Ω if it were taken as a stand-alone portfolio, i.e.,

$$\widetilde{\mathrm{GA}}_\Omega = \frac{1}{2\mathcal{K}^*_\Omega} \sum_{i=1}^{n} \left(\frac{s_i}{S_\Omega}\right)^2 Q_i \mathcal{C}_i.$$

It is straightforward to decompose the whole portfolio GA as

$$\begin{aligned}\widetilde{\mathrm{GA}} &= \frac{1}{2\mathcal{K}^*} \sum_{i=1}^{n} s_i^2 Q_i \mathcal{C}_i = \frac{1}{2\mathcal{K}^*} \sum_{i\in\Omega} s_i^2 Q_i \mathcal{C}_i + \frac{1}{2\mathcal{K}^*} \sum_{i\in\Omega^c} s_i^2 Q_i \mathcal{C}_i \\ &= S_\Omega^2 \frac{\mathcal{K}^*_\Omega}{\mathcal{K}^*} \widetilde{\mathrm{GA}}_\Omega + (1 - S_\Omega)^2 \frac{\mathcal{K}^*_{\Omega^c}}{\mathcal{K}^*} \widetilde{\mathrm{GA}}_{\Omega^c}. \qquad (22)\end{aligned}$$

This leads directly to the lower bound,

$$\widetilde{\mathrm{GA}} \geq S_\Omega^2 \frac{\mathcal{K}^*_\Omega}{\mathcal{K}^*} \widetilde{\mathrm{GA}}_\Omega. \qquad (23)$$

This bound can be useful when Ω represents the wholesale banking book and Ω^c the retail portfolio. The high name diversification of retail portfolios suggests that $\widetilde{\mathrm{GA}}_{\Omega^c}$ may be negligibly small, in which case the bound of equation (23) will be tight.

4. Empirical Magnitude and Sensitivity Analysis

In this section, we employ loan data from the German credit register to study the empirical behavior and comparative statics of the granularity adjustment. The data are described in section 4.1. Numerical results are reported in section 4.2. Finally, in section 4.3, we use the data to calibrate the precision parameter ξ.

4.1 Data on German Bank Portfolios

We construct test portfolios from two Bundesbank data sets. First, the German credit register is a quarterly census of loans issued by banking institutions in Germany. At each quarter-end, banks are required to report all loans and other credit exposures to borrowers above an aggregate exposure threshold of 1.5 million euros. Second, the Bundesbank collects annual balance sheet data on a large sample of German firms. Gerke et al. (2008) model firm annual PD as a function of these balance sheet data by fitting a logistic regression model to default by firms in this data set. We draw on their results to estimate PD for borrowers in our sample.[8]

In terms of exposure count, our constructed portfolios may be much smaller than the banks' actual portfolios. This is due in part to the threshold imposed in the credit register. Schmieder (2006, §3.2) reports that the credit register covers 20 percent of household debt, 60 percent of corporate debt, and nearly 100 percent of interbank debt in Germany. Altogether, roughly 70 percent of all German borrowing is captured, but clearly the smallest borrowers are underrepresented.[9] Compounding this selection problem is that small and medium enterprises (SMEs) are underrepresented in the balance sheet data (Gerke et al. 2008). An implication is that the granularity adjustments for our constructed portfolios will tend to overstate the GA for complete portfolios of comparable quality. It is reasonable to assume that the sub-portfolio of missing borrowers is a retail portfolio with negligible granularity adjustment on a stand-alone basis. If the constructed portfolio represents 70 percent of the whole portfolio by exposure, then the analysis of section 3.3 suggests that our stand-alone GA for the constructed portfolio is roughly twice as large as the GA would be for the whole portfolio.

Our data contain portfolios for thirty-eight German banks as of end-2002 with at least 250 exposures in the matched sample. We group the banks into categories labeled Large, Medium, Small, and Very Small. The large banks have at least 4,000 exposures; medium banks have 1,000–3,999 exposures; small banks have 500–999 exposures; and very small banks have 250–499 exposures. These groups contain four, six, eleven, and seventeen banks, respectively.

As a benchmark, we form a reference portfolio based on aggregate data. To accommodate privacy restrictions on credit register

[8] We thank Christian Schmieder for providing the fitted PDs.

[9] The selection bias is mitigated by the way the minimum threshold is applied. So long as a loan exceeds 1.5 million euros at some date in the quarter (not necessarily at the end of the quarter), the loan must be reported. Also, the threshold is applied to related groups of borrowers, so loans to individual borrowers can be smaller.

Table 1. Exposure Distribution in Portfolio $\mathcal{P}$

Level	Quantile
5%	50.92
25%	828.80
Median	1,811.75
75%	3,705.50
95%	13,637.36

Note: This table reports quantiles of the loan size distribution in our reference portfolio $\mathcal{P}$ in thousands of euros.

data, we merge portfolios for three large banks into a single data set. We sort the merged portfolio by exposure size and retain only every fourth exposure. The resulting portfolio of 5,289 borrowers is somewhat smaller than the mean large portfolio but realistic in terms of exposure and PD distribution. The mean of the loan size distribution is 3,973 euros and the standard deviation is 9,435 euros. Quantiles are reported in table 1. Henceforth, this reference portfolio is denoted as portfolio $\mathcal{P}$.

We associate ranges of PD values to discrete rating grades, for which we adopt S&P whole-letter rating-grade nomenclature. In table 2, we report the PD ranges and the distribution of borrowers in portfolio $\mathcal{P}$ across rating grades. The mean PD of portfolio $\mathcal{P}$ is 43 basis points, and the exposure-weighted mean is 50 basis points. Including all matched loans in our sample, the mean PD is 44 basis points, which suggests that portfolio $\mathcal{P}$ is representative of the entire sample in terms of credit quality. However, Escott, Glormann, and Kocagil (2011, p. 8) find that the mean PD within small- and medium-sized German loan portfolios is roughly 80 basis points. The difference is probably attributable to a bias in the balance sheet sample towards higher credit quality firms, particularly for the SME sub-sample (Gerke et al. 2008, §3.ii). As a consequence, the GA for our portfolio $\mathcal{P}$ is likely to understate the GA for an actual bank portfolio of similar exposure distribution.

Our data do not contain information on LGD or loan maturity, so we impose the foundation IRB assumptions of E[LGD] $= 0.45$ and effective maturity of $M = 2.5$ years.

Table 2. Borrower Distribution by Rating Grade

Grade	PD Range (%)	Frequency (%)	
		Unweighted	**Exposure Weighted**
AAA	≤ 0.02	2.25	0.90
AA	0.02–0.06	6.07	3.63
A	0.06–0.18	23.80	18.82
BBB	0.18–1.06	60.45	66.41
BB	1.06–4.94	7.24	10.04
B	4.94–19.14	0.19	0.20
CCC	> 19.14	0.00	0.00

Notes: This table displays our mapping of borrower PD ranges into S&P whole-letter rating-grade nomenclature. The third and fourth columns report, respectively, the unweighted and exposure-weighted distribution of borrowers in reference portfolio $\mathcal{P}$.

4.2 Magnitude and Sensitivity Analysis

For each bank portfolio in our sample, we calculate the exact GA of formula (14). The results are grouped by bank size category and reported in table 3. Here, and throughout this section, we fix a quantile level of $q = 0.999$ and assume that the precision parameter $\xi = 0.125$, which implies $\delta = 4.3$.

In the fourth column, the GA is expressed as a percentage of total exposure. We see that the GA is almost negligible (12 to 13 basis points) for the large portfolios, but can be material (up to 96 basis points) for the medium portfolios and can be quite substantial for small and (especially) very small portfolios. In the final column, the GA is reported as a percentage of unexpected loss (approximated as $\mathcal{K}$ + GA), which provides a measure of the relative impact on required capital. In a study based on applying an actuarial multifactor CreditMetrics model to U.S. syndicated loan portfolio data, Heitfield, Burton, and Chomsisengphet (2006) find that name concentration accounts for between 1 percent and 8 percent of VaR depending on the portfolio size. Except for one outlier among the medium portfolios, the GA-to-VaR ratio for our large and medium portfolios fall within the same range. Tarashev and Zhu (2008) find comparable results as well, and indeed confirm that our CreditRisk$^+$ GA formula is highly accurate as an approximation to the difference between actual (i.e., simulated) VaR and asymptotic VaR for their

Table 3. Granularity Adjustment for German Bank Portfolios

Portfolio	Number of Exposures	HHI	Granularity Adjustment	GA/UL
$\mathcal{P}^*$	5,289	0.0002	0.02	0.37
$\mathcal{P}$	5,289	0.0013	0.15	3.07
Large	4,000–8,999	≤ 0.0012	0.12–0.13	2.37–2.71
Medium	1,000–3,999	0.001–0.011	0.14–0.96	3.17–17.27
Small	500–999	0.004–0.011	0.36–1.14	6.57–18.65
Very Small	250–499	0.005–0.031	0.48–3.81	7.99–40.18

Notes: This table reports the GA for the reference portfolio $\mathcal{P}$ and for thirty-eight German bank portfolios with at least 250 exposures in the matched data sample. The bank portfolios were sorted by exposure count into categories labeled Large, Medium, Small, and Very Small, containing four, six, eleven, and seventeen banks, respectively. Portfolio $\mathcal{P}^*$ is a homogenized analog to reference portfolio $\mathcal{P}$ with $n = 5{,}289$ equal-sized loans and constant PD = 0.43 percent. The third column reports the range of the Herfindahl-Hirschman Index within each set of banks. The fourth column reports the corresponding range of GA calculated using equation (14), which is expressed as a percentage of total exposure. The final column reports the range of the ratio of GA to unexpected loss.

two-state CreditMetrics model.[10] As Tarashev and Zhu (2008) argue, the magnitude of the GA should be compared against other potential sources of error in capital assessment. By simulating data sets of plausible length and estimating asset correlation parameters on the simulated data, they show that the uncertainty in VaR due to estimation error in the level and dispersion of correlation is plausibly 18–30 percent of VaR, which is significantly larger than the GA for portfolios of medium size and larger.

The GA for our reference portfolio $\mathcal{P}$ places it at the border between the medium and large bank portfolios. For comparison against $\mathcal{P}$, we create a *homogeneous* reference portfolio with the same number of loans ($n = 5{,}289$). In this portfolio $\mathcal{P}^*$, the loans are of equal size and each borrower has PD of 43 basis points. We find that the GA for $\mathcal{P}^*$ is only one-eighth as large as the GA for $\mathcal{P}$.

[10]For the same two-state CreditMetrics model, Gürtler, Heithecker, and Hibbeln (2008) demonstrate that the GA formula of Emmer and Tasche (2005) is highly accurate in portfolios of very modest size.

Table 4. Approximation Error of the Simplified $\widetilde{\mathrm{GA}}$

Portfolio	P0	P1	P2	P10	P50
PD = 1%					
Exposure A_i	1	i	i^2	i^{10}	i^{50}
$\widetilde{\mathrm{GA}}$	10.48	13.97	18.86	60.36	269.71
GA	10.79	14.38	19.41	62.13	277.62
PD = 4%					
Exposure A_i	1	i	i^2	i^{10}	i^{50}
$\widetilde{\mathrm{GA}}$	11.75	15.66	21.14	67.66	302.35
GA	12.34	16.45	22.21	71.08	317.64

Notes: This table demonstrates the accuracy of the simplified $\widetilde{\mathrm{GA}}$ given by equation (17) as an approximation to the exact GA of equation (14). We calculate exact and simplified GA for stylized portfolios with different degrees of exposure concentrations. Portfolio P0 is a homogeneous portfolio of $n = 1{,}000$ loans where all loans are of size 1. Exposure A_i for loan $i = 1, \ldots, 1000$ in portfolio P1 (respectively, P2, P10, P50) is given as i (respectively, i^2, i^{10}, i^{50}). Borrower PD is fixed at 1 percent in the upper panel of the table and at 4 percent in the lower panel. We fix E[LGD] at 45 percent, maturity at $M = 2.5$ years, quantile level to $q = 0.999$, and the precision parameter to $\xi = 0.125$. GA is expressed in basis points.

This illustrates the importance of portfolio heterogeneity in credit concentrations.

The third column of table 3 reports the range of the Herfindahl Index for each group of banks. The results demonstrate the strong correlation between the HHI and GA across these portfolios. The correspondence is not exact, as the GA is sensitive to credit quality as well.

Our next task is to assess the accuracy of the simplified granularity adjustment $\widetilde{\mathrm{GA}}$ as an approximation to the exact GA. We construct six stylized portfolios differing in exposure concentration. Each portfolio consists of $n = 1{,}000$ exposures and has constant PD. Portfolio P0 is completely homogeneous. In portfolio Pi, loan j has exposure size $A_j = j^i$. In our most concentrated portfolio (P50), the largest exposure accounts for 5 percent of the total exposure. The values for both the simplified $\widetilde{\mathrm{GA}}$ and the exact GA for each of these portfolios are listed in table 4. We see that the approximation error increases with concentration and with PD. For realistic portfolios, the error is trivial. Even for the case of portfolio P10 and PD = 4

Figure 1. Tightness of the Upper Bound

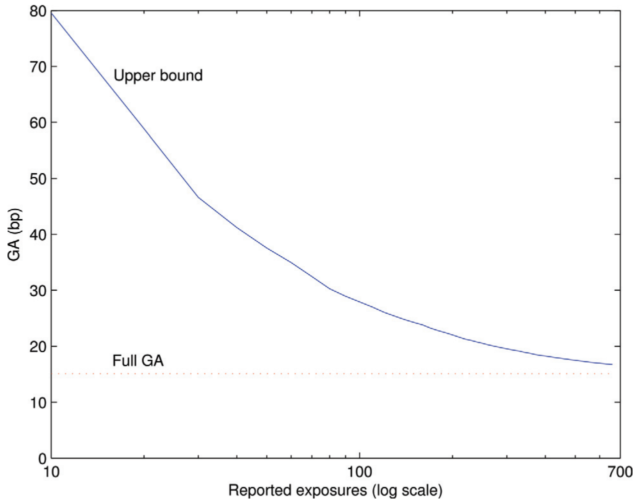

Notes: This figure shows the convergence of the upper bound on the GA to the GA as the number of loans included in the reported set Ω converges to the number of loans (n) in the whole portfolio. The $\widetilde{\mathrm{GA}}$ for the whole portfolio (horizontal line) is calculated using the simplified formula (17) while the upper bound is computed by equation (21). The portfolio is our reference portfolio $\mathcal{P}$. We fix E[LGD] at 45 percent, maturity $M = 2.5$ years, quantile level $q = 0.999$, and the precision parameter $\xi = 0.125$. The GA and its upper bound are expressed in basis points.

percent, the error is only 3 basis points. The error grows to 16 basis points in the extreme example of P50 and PD = 4 percent, but even this remains small relative to the size of the GA.

Finally, we demonstrate the effectiveness of the upper bound provided in section 3 on portfolio $\mathcal{P}$. In figure 1, we show how the gap between the upper bound and the "whole portfolio" $\widetilde{\mathrm{GA}}$ shrinks as the number of included positions increases. With only $m = 140$ exposures included out of $n = 5{,}289$ in the whole portfolio, this gap is under 10 basis points. With $m = 270$ exposures included, the gap shrinks to 5 basis points.

4.3 Calibration of the Precision Parameter

We have so far taken the CreditRisk$^+$ precision parameter ξ as known. In this section, we recognize that ξ is not known and therefore needs to be calibrated. In principle, ξ could be estimated by maximum likelihood given panel data on default frequencies by rating grade. When factor loadings are also free parameters to be estimated, ξ is identified through its effect on the skewness and higher moments of the distribution of default frequencies (see Gordy 2000, §4.2), and therefore cannot be precisely estimated given the short time series that are available.[11] Since our main objective is practical regulatory application, we propose instead to calibrate ξ to be as consistent as possible with the IRB model. Specifically, we find the value of ξ that equates the $\widetilde{\mathrm{GA}}$ of equation (17) to the GA implied by a close proxy to the IRB model for a representative portfolio. As noted in section 2, the model underpinning the IRB formula is a hybrid of industry models with Gaussian factor structure. As an approximation to this model, we use a mark-to-market CreditMetrics model.[12]

CreditMetrics and other Gaussian models have no parameter that corresponds directly to ξ in CreditRisk$^+$. When the systematic factor is normally distributed, the shape of the distribution for X does not depend on its variance. For this reason, one can normalize the variance to 1 without any loss of generality. By contrast, when X is gamma-distributed as in CreditRisk$^+$, skewness and kurtosis and other shape measures for X are not invariant to the variance, and so this parameter cannot be normalized. Higher values of ξ imply lower systematic risk, which generally leads to lower economic capital requirements, but which maximizes the GA as a share of economic capital. The economic significance of the effect on capital is demonstrated by Gordy (2000, table 6).

Recall that ξ influences the GA through the δ parameter. In figure 2, we show that δ is nearly log-linear in ξ across the plausible range of values for ξ and tail quantile levels q. Because credit loss distributions are highly skewed, capital contributions ($\mathcal{K}$) are

[11] In industry practice, we believe it is most common either to impose an arbitrary round value such as $\xi = 1$ or to fix the factor loading as $\omega = 1$ and estimate ξ using the variance of the default frequency.

[12] The original version of this model is documented by Gupton, Finger, and Bhatia (1997). The current MSCI RiskMetrics implementation adopts the Hull and White (2000) pricing model for repricing positions at the horizon. As this allows for greater internal consistency than the original discount cashflow pricing rules, we follow suit.

Figure 2. Effect of Precision ξ on δ

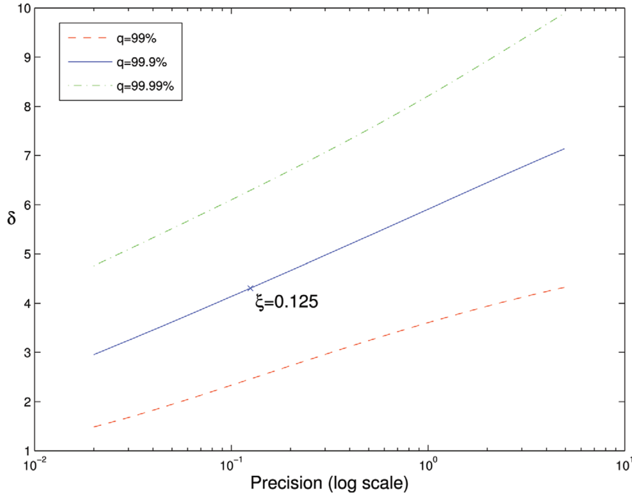

Notes: This figure depicts the effect of parameter ξ on the regulatory parameter δ defined in equation (15). ξ controls the precision of the systematic risk factor X in the CreditRisk$^+$ model (i.e., $V[X] = 1/\xi$), and is plotted on the x-axis in log scale. The solid line is for the baseline quantile level $q = 99.9$ percent. The dashed line for $q = 99$ percent and dash-dotted panel for $q = 99.99$ percent demonstrate the robustness of the log-linear relationship. The point marked with × shows that $\delta = 4.31$ for baseline values $\xi = 0.125$ and $q = 0.999$ used in section 4.2.

much larger than loan-loss reserve contributions ($\mathcal{R}$), so $\widetilde{\mathrm{GA}}$ should be roughly proportional to $\delta - 1$. This is indeed what we find. In figure 3, we plot $\widetilde{\mathrm{GA}}$ for portfolio $\mathcal{P}$ against δ. When the assumed value for δ is increased from 3 to 5, the GA very nearly doubles.[13]

Our baseline experiment is constructed as follows. Our portfolio is the German reference portfolio $\mathcal{P}$. Each borrower is mapped to an S&P rating grade using the PD ranges given in table 2. We assume that each loan is a newly issued term loan of 2.5 years, and that E[LGD] is 45 percent. We assign to each loan the par coupon given by the Hull and White (2000) pricing model calibrated to ratings-based spread indices on December 31, 2004, when prevailing spreads

[13]Under our baseline assumptions of $q = 0.999$, $M = 2.5$ years, and LGD $= 0.45$, $\widetilde{\mathrm{GA}}$ for portfolio $\mathcal{P}$ is 9.18 basis points at $\delta = 3$ and 18.19 basis points at $\delta = 5$.

Figure 3. Effect of δ on $\widetilde{\mathrm{GA}}$

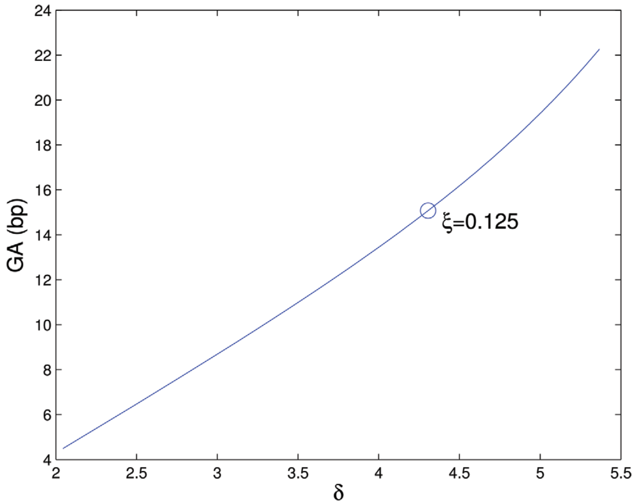

Notes: This figure depicts the effect of the parameter δ defined in equation (15) on the simplified $\widetilde{\mathrm{GA}}$ of equation (17) for reference portfolio $\mathcal{P}$. We fix our baseline assumptions of quantile level $q = 0.999$, maturity $M = 2.5$ years, and E[LGD] $= 0.45$. The point corresponding to baseline assumption $\xi = 0.125$ is marked. The GA is expressed in basis points.

were roughly at the midpoint for the 2003–07 period.[14] We calculate effective maturity M using the duration formula specified by Basel II (Basel Committee on Banking Supervision 2006, §320). The matrix of unconditional transition probabilities is taken from Vazza, Aurora, and Kraemer (2009, table 13).[15] The asset-correlation parameter is determined at the grade level by the IRB formula (Basel Committee on Banking Supervision 2006, §272), which assigns asset correlation as a function of the grade's associated default probability.

[14]We thank Chris Finger of MSCI RiskMetrics for providing these spreads and the details of the current pricing approach in CreditMetrics. The swap spread on the CDX North America investment-grade five-year index was 43.8 basis points on December 31, 2004.

[15]The published matrix includes transitions to "not rated." We remove NR as an end state and adjust the other probabilities in each row of the matrix by dividing by $1 - \Pr(\mathrm{NR})$ for that row.

Table 5. Calibrated Values of the Precision Parameter

Portfolio		Maturity		
		1 Year	**2.5 Years**	**5 Years**
$\mathcal{P}$	δ	3.6488	3.4668	3.2272
	ξ	0.0520	0.0407	0.0292
BBB	δ	3.5287	3.2498	2.9890
	ξ	0.0442	0.0302	0.0209
BB	δ	4.1779	4.0570	3.8994
	ξ	0.1056	0.0899	0.0729

Notes: This table displays the calibrated values of the CreditRisk$^+$ precision parameter ξ and the regulatory parameter δ defined by equation (15). Calibration is performed by matching the GA implied by CreditMetrics to the simplified $\widetilde{\text{GA}}$ of equation (17). The panel shows results for the reference portfolio $\mathcal{P}$ and for homogeneous portfolios of grade BBB and BB for maturities of 1 year, 2.5 years, and 5 years.

We choose δ so that the simplified GA of equation (17) matches the GA implied by CreditMetrics. To implement equation (17), we calculate for each loan $\mathcal{K}_i$ and $\mathcal{R}_i$ by the IRB formulae. The CreditMetrics GA is calculated analytically by the method of Gordy and Marrone (2012). Results are displayed in table 5. For the baseline case of portfolio $\mathcal{P}$ with 2.5 years maturity, we find $\delta = 3.47$ ($\xi = 0.04$). As robustness checks, we also obtain δ for homogeneous portfolios of grade BBB and grade BB, and for maturities of one and five years. We find higher δ for the lower-rated portfolio, and also that δ decreases with maturity.[16] As an additional robustness check, we calibrated CreditMetrics using the spreads of December 29, 2006, when market spreads were at a cyclical low point, and obtain very similar results. We conclude that fixing $\delta = 4$ would be parsimonious and consistent with our results. However, the baseline value of $\xi = 0.125$ ($\delta = 4.3$) used above in section 4.2 could be justified as an appropriately conservative choice for a typical bank portfolio.

[16]We have calculated δ for homogeneous portfolios of each rating grade (not reported) and find that δ varies monotonically with rating (lowest for AAA, highest for CCC).

While the absolute magnitude of the GA is sensitive to δ, its relative magnitude across bank portfolios is much less so. When we sort our thirty-eight German bank portfolios by $\widetilde{\mathrm{GA}}$, the ordering depends minimally on δ. Decreasing δ from 5 to 3 reverses the order between two very small banks but otherwise leaves the ordering unchanged. Increasing δ from 5 to 7 does not alter the ordering at all.[17] Thus, the proper functioning of the GA as a supervisory tool does not materially depend on the precision with which δ is calibrated.

5. Conclusion

This paper sets forth a granularity adjustment for portfolio credit VaR that accounts for undiversified idiosyncratic risk in the portfolio. This source of risk is not captured by the minimum capital requirement of the IRB approach of Basel II and III. For large bank portfolios, the GA is generally negligible, but for the small and medium-sized portfolios in our data set of German bank loans, the GA ranges between 3 percent and 20 percent of total capital.

Among the modifications we have made to earlier proposals for granularity adjustment, most important is the introduction of an upper-bound methodology that allows calculation of the GA based on incomplete information. This addresses the most significant source of operational burden associated with implementation. For some banks, this approach could permit significant reductions in data requirements at modest cost in additional capital requirement.

We note two limitations of our methodology. First, the GA formula is itself an asymptotic approximation, and so might not work well on very small portfolios. We do not see this issue as a material concern. In general, the GA errs on the conservative side (i.e., it overstates the effect of granularity) but is quite accurate for modest-sized portfolios of as few as 200 borrowers (for a low-quality portfolio) or 500 borrowers (for an investment-grade portfolio). Second, the IRB formulae are based on a rather different model of credit risk, so we have a form of "basis risk" in the model mismatch. This is potentially

[17] Let Y_i^δ be the $\widetilde{\mathrm{GA}}$ for bank i for a given δ. Kendall's tau between vectors Y^3 and Y^5 and between vectors Y^3 and Y^7 is 0.9972.

a more serious issue but is mitigated by the way we parameterize our GA and calibrate its underlying model. By parameterizing the GA formula in terms of the loan-level reserve requirements and IRB capital charges, we implicitly force the GA model into agreement with the IRB model on both expected loss and asymptotic value-at-risk. Furthermore, we calibrate a free parameter in our model so that the GA matches simulation-based estimates of the "true" GA for the IRB model on a representative portfolio.

The great advantage to the particular model we use to underpin the GA is its analytical tractability. The linear specification of conditional default probability that is peculiar to this model is what makes it possible to parameterize the GA formula in terms of the IRB "outputs" (i.e., loan-loss reserve requirements and capital charges) in place of the model primitives. As these outputs must be calculated for Basel compliance, our GA can be implemented without imposing additional data requirements. Finally, without the analytical tractability of our approach, it would not have been possible to derive a useful upper-bound methodology.

Appendix. Granularity Adjustment for Expected Shortfall

In the case of a (locally) continuous loss Y, *expected shortfall* (ES) is the expected loss conditional on exceeding a VaR threshold. Formally,

$$\mathrm{ES}_q[Y] = \mathrm{E}[Y|Y \geq \alpha_q(Y)] = \frac{1}{1-q}\int_q^1 \alpha_u(Y)\, du. \tag{24}$$

Acerbi and Tasche (2002) show that ES is coherent in the sense of Artzner et al. (1999) and equivalent to the "conditional VaR" (CVaR) measure of Rockafellar and Uryasev (2002). The granularity adjustment for ES is given by Gordy (2004) and Martin and Tasche (2007) as

$$\mathrm{GA}^{ES} = \frac{1}{2(1-q)} h(\alpha_q(X)) \frac{\sigma^2(\alpha_q(X))}{\mu'(\alpha_q(X))}. \tag{25}$$

It is readily checked that the GA for VaR of equation (9) can be obtained as the derivative of GA^{ES} with respect to q (times a sign change). Martin and Tasche (2007) show that GA^{ES} is always positive for q near 1, whereas the VaR GA can be negative in certain pathological situations.[18]

Proceeding as in section 2, the exact granularity adjustment for CreditRisk$^+$ under expected shortfall is

$$\mathrm{GA}^{ES} = \frac{1}{2\mathcal{K}^*}\Delta \sum_{i=1}^{N} s_i^2 \left(\mathcal{C}_i(\mathcal{K}_i + \mathcal{R}_i) + (\mathcal{K}_i + \mathcal{R}_i)^2 \frac{\mathrm{V}[\mathrm{LGD}_i]}{\mathrm{E}[\mathrm{LGD}_i]^2} \right), \tag{26}$$

where

$$\begin{aligned} \Delta &\equiv \frac{(\alpha_q(X) - 1)h(x_q)}{1-q} \\ &= \frac{(\alpha_q(X) - 1)}{1-q} \frac{\xi^\xi}{\Gamma(\xi)} \alpha_q(X)^{\xi-1} \exp(-\xi\alpha_q(X)). \end{aligned}$$

The regulatory parameter Δ takes the place of δ in the VaR GA for CreditRisk$^+$. It depends only on model parameters q and ξ, and it is through Δ that ξ influences the ES GA. When $\xi = 0.25$ and $q = 0.999$, we have $\Delta = 4.73$ (compared with $\delta = 4.83$ for the VaR GA), which implies that the ES GA is similar in magnitude to the VaR GA. Ignoring second-order terms in $\mathcal{K}_i$ and $\mathcal{R}_i$, we can simplify the ES GA as

$$\widetilde{\mathrm{GA}}^{ES} = \frac{1}{2\mathcal{K}^*}\Delta \sum_{i=1}^{N} s_i^2 \mathcal{C}_i(\mathcal{K}_i + \mathcal{R}_i). \tag{27}$$

This is nearly identical in form to the VaR-simplified GA of equation (17).

Finally, we can bound the ES-simplified GA along the lines of section 3. Following the notation of section 3.2, we find that $\widetilde{\mathrm{GA}}^{ES}$ is bounded above by the following:

[18]The VaR GA can be negative when there is a local mode in the density of the loss distribution in the neighborhood of the asymptotic VaR.

$$\frac{1}{2\mathcal{K}^*}\Delta\left(\sum_{i\in\Omega} s_i^2\mathcal{C}_i(\mathcal{K}_i+\mathcal{R}_i)+\bar{s}\left((\mathcal{K}^*-S_\Omega\mathcal{K}_\Omega^*)+(\mathcal{R}^*-S_\Omega\mathcal{R}_\Omega^*)\right)\right).$$

The lower bound takes the same form as the lower bound on the VaR-simplified GA in equation (23), i.e.,

$$\widetilde{\mathrm{GA}}^{ES} \geq S_\Omega^2\frac{\mathcal{K}_\Omega^*}{\mathcal{K}^*}\widetilde{\mathrm{GA}}_\Omega^{ES},$$

where $\widetilde{\mathrm{GA}}_\Omega^{ES}$ is the ES-simplified GA for Ω as a stand-alone portfolio.

References

Acerbi, C., and D. Tasche. 2002. "On the Coherence of Expected Shortfall." *Journal of Banking and Finance* 26 (7): 1487–1503.

Acharya, V. V., S. T. Bharath, and A. Srinivasan. 2007. "Does Industry-Wide Distress Affect Defaulted Firms? Evidence from Creditor Recoveries." *Journal of Financial Economics* 85 (3): 787–821.

Altman, E. I., B. Brady, A. Resti, and A. Sironi. 2005. "The Link between Default and Recovery Rates: Theory, Empirical Evidence, and Implications." *Journal of Business* 78 (6): 2203–28.

Antonov, A., S. Mechkov, and T. Misirpashaev. 2007. "Analytical Techniques for Synthetic CDOs and Credit Default Risk Measures in Static Factor Models." *Wilmott Magazine* 7 (November–December): 84–105.

Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath. 1999. "Coherent Measures of Risk." *Mathematical Finance* 9 (3): 203–28.

Basel Committee on Banking Supervision. 2001. "Basel II: The New Basel Capital Accord." Second Consultative Paper, Bank for International Settlements (January).

———. 2006. "Basel II: International Convergence of Capital Measurement and Capital Standards: A Revised Framework." Publication No. 128, Bank for International Settlements (June).

———. 2011. "Basel III: A Global Regulatory Framework for More Resilient Banks and Banking Systems." Publication No. 189, Bank for International Settlements (June).

Counterparty Risk Management Policy Group III. 2008. "Containing Systemic Risk: The Road to Reform." Technical Report (August).

Credit Suisse Financial Products. 1997. "CreditRisk+: A Credit Risk Management Framework."

Ebert, S., and E. Lütkebohmert. 2011. "Treatment of Double Default Effects within the Granularity Adjustment for Basel II." *Journal of Credit Risk* 7 (1): 3–33.

Emmer, S., and D. Tasche. 2005. "Calculating Credit Risk Capital Charges with the One-Factor Model." *Journal of Risk* 7 (2): 85–101.

Escott, P., F. Glormann, and A. E. Kocagil. 2011. "RiskCalc for Private Companies: The German Model." Modeling Methodology, Moody's KMV (November).

Gagliardini, P., and C. Gouriéroux. 2011. "Approximate Derivative Pricing for Large Classes of Homogeneous Assets with Systematic Risk." *Journal of Financial Econometrics* 9 (2): 237–80.

Gagliardini, P., C. Gouriéroux, and A. Monfort. 2012. "Microinformation, Nonlinear Filtering, and Granularity." *Journal of Financial Econometrics* 10 (1): 1–53.

Garcia Cespedes, J. C., J. A. 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.

Gerke, W., F. Mager, T. Reinschmidt, and C. Schmieder. 2008. "Empirical Risk Analysis of Pension Insurance: The Case of Germany." *Journal of Risk and Insurance* 75 (3): 763–84.

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.

———. 2004. "Granularity Adjustment in Portfolio Credit Risk Measurement." In *Risk Measures for the 21st Century*, ed. G. P. Szegö, 109–22 (chapter 8). John Wiley & Sons.

Gordy, M. B., and S. Juneja. 2010. "Nested Simulation in Portfolio Risk Measurement." *Management Science* 56 (10): 1833–48.

Gordy, M. B., and J. Marrone. 2012. "Granularity Adjustment for Mark-to-Market Credit Risk Models." *Journal of Banking and Finance* 36 (7): 1896–1910.

Gouriéroux, C., and J. Jasiak. 2012. "Granularity Adjustment for Default Risk Factor Model with Cohorts." *Journal of Banking and Finance* 36 (5): 1464–77.

Gouriéroux, C., J.-P. Laurent, and O. Scaillet. 2000. "Sensitivity Analysis of Values at Risk." *Journal of Empirical Finance* 7 (3–4): 225–45.

Gouriéroux, C., and A. Monfort. 2009. "Granularity in a Qualitative Factor Model." *Journal of Credit Risk* 5 (4): 29–61.

Gundlach, V. M., and F. B. Lehrbass, eds. 2004. *CreditRisk+ in the Banking Industry.* Heidelberg: Springer-Verlag.

Gupton, G. M., C. C. Finger, and M. Bhatia. 1997. *CreditMetrics—Technical Document.* New York: J. P. Morgan & Co. (April).

Gürtler, M., D. Heithecker, and M. Hibbeln. 2008. "Concentration Risk under Pillar 2: When Are Credit Portfolios Infinitely Fine Grained?" *Kredit und Kapital* 41 (1): 79–124.

Gürtler, M., M. Hibbeln, and C. Vöhringer. 2010. "Measuring Concentration Risk for Regulatory Purposes." *Journal of Risk* 12 (3): 69–104.

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. 2000. "Valuing Credit Default Swaps I: No Counterparty Default Risk." *Journal of Derivatives* 8 (1): 29–40.

Institute of International Finance. 2009. "Reform in the Financial Services Industry: Strengthening Practices for a More Stable System." Report of the IIF Steering Committee on Implementation (December).

Lütkebohmert, E. 2012. "Failure of the Saddlepoint Method in the Presence of Double Defaults." *Journal of Risk* 15 (1): 71–89.

Martin, R., and D. Tasche. 2007. "Shortfall: A Tail of Two Parts." *Risk* 20 (2): 84–89.

Martin, R., and T. Wilde. 2002. "Unsystematic Credit Risk." *Risk* 15 (11): 123–28.

McNeil, A. J., R. Frey, and P. Embrechts. 2005. *Quantitative Risk Management: Concepts, Techniques and Tools.* Princeton, NJ: Princeton University Press.

Merino, S., and M. Nyfeler. 2004. "Numerical Techniques for Determining Portfolio Credit Risk." In *CreditRisk+ in the Banking Industry*, ed. V. M. Gundlach and F. B. Lehrbass, 279–309 (chapter 17). Heidelberg: Springer-Verlag.

Panjer, H. H., and G. E. Willmot. 1992. *Insurance Risk Models.* Schaumberg, IL: Society of Actuaries.

Pykhtin, M. 2004. "Multi-Factor Adjustment." *Risk* 17 (3): 85–90.

Rockafellar, R. T., and S. Uryasev. 2002. "Conditional Value-at-Risk for General Loss Distributions." *Journal of Banking and Finance* 26 (7): 1443–71.

Schmieder, C. 2006. "The Deutsche Bundesbank's Large Credit Database (BAKIS-M and MiMiK)." *Schmollers Jahrbuch* 126 (4): 653–63.

Tarashev, N., C. Borio, and K. Tsatsaronis. 2010. "Attributing Systemic Risk to Individual Institutions." BIS Working Paper No. 308 (May).

Tarashev, N., and H. Zhu. 2008. "Specification and Calibration Errors in Measures of Portfolio Credit Risk: The Case of the ASRF Model." *International Journal of Central Banking* 4 (2): 129–73.

Vasicek, O. A. 2002. "Loan Portfolio Value." *Risk* 15 (12): 160–62.

Vazza, D., D. Aurora, and N. Kraemer. 2009. "2008 Annual Global Corporate Default Study and Rating Transitions." Technical Report, Standard & Poor's (April).

Voropaev, M. 2011. "An Analytical Framework for Credit Portfolio Risk Measures." *Risk* 24 (5): 72–78.

Wilde, T. 2001. "Probing Granularity." *Risk* 14 (8): 103–6.