A Risk-Factor Model Foundation for Ratings-Based Bank Capital Rules

Michael B. Gordy*
Board of Governors of the Federal Reserve System

Draft: February 5, 2001

Abstract

When economic capital is calculated using a portfolio model of credit value-at-risk, the marginal capital requirement for an instrument depends, in general, not only on the characteristics of the instrument itself, but also on the properties of the portfolio in which it is held. By contrast, ratings-based risk-bucket capital rules, including both the current Basel Accord and any its recently proposed replacement, assign a capital charge to an instrument based only on its own characteristics.

In this paper, I demonstrate that risk-bucket capital rules can be reconciled with the general class of credit VaR models. VaR models imply marginal capital charges which depend only on an asset's own characteristics only if (a) there is only a single systematic risk factor driving correlations across obligors, and (b) portfolios are "asymptotically fine-grained," i.e., no exposure in a portfolio accounts for more than an arbitrarily small share of total exposure. This result holds under very general assumptions on the distribution of exposure sizes and credit ratings in the portfolio. It allows for variation across obligors in sensitivity to the systematic risk factor; for systematic risk in recovery rates; and for either actuarial or mark-to-market notions of loss.

Analysis of rates of convergence to asymptotic VaR is used to design a simple method of approximating a portfolio-level add-on charge for undiversified idiosyncratic risk. Thus, violation of the second assumption need not pose a practical problem. There is no similarly simple way to address violation of the single factor assumption. Regulators ought therefore to be especially cautious in applying risk-bucket capital rules to portfolios with significant industry and geographic concentrations. Although unlikely to pose a near-term obstacle to Basel reform, dependence on this assumption may limit the long-term viability of risk-bucket rules for regulatory capital.

*The views expressed herein are my own and do not necessarily reflect those of the Board of Governors or its staff. I would like to thank Paul Calem, Darrell Duffie, Jon Faust, David Jones and David Lando for their helpful comments, and Susan Yeh for editorial suggestions. Please address correspondence to the author at Division of Research and Statistics, Mail Stop 153, Federal Reserve Board, Washington, DC 20551, USA. Phone: (202)452-3705. Fax: (202)452-5295. Email: ⟨michael.gordy@frb.gov⟩.

Recent years have witnessed significant advances in the design, calibration and implementation of portfolio models of credit risk. Large commercial banks and other financial institutions with significant credit exposure rely increasingly on models to guide credit risk management at the portfolio level. Models allow management to identify concentrations of risk and opportunities for diversification within a disciplined and objective framework, and thus offer a more sophisticated, less arbitrary alternative to traditional lending limit controls. More widespread and intensive use of models is encouraging a more active approach to portfolio management at commercial banks, which has contributed to the improved liquidity of markets for debt instruments and credit derivatives.

Stripped to its essentials, a credit risk model is a function mapping from a parsimonious set of instrument-level characteristics and market-level parameters to a distribution for portfolio credit losses over some chosen horizon. The model output of primary interest, the "economic capital" required to support the portfolio, is derived as some summary statistic of the loss distribution. The definition of economic capital in most widespread use is value-at-risk ("VaR"). Under the VaR paradigm, an institution holds capital in order to maintain a target rating for its own debt. Associated with the target rating is a probability of survival over the horizon (say, 99.9% over one year). To be consistent with its target survival probability (denoted q), the institution must hold reserves and equity capital sufficient to cover up to the q^{th} quantile of the distribution of portfolio loss over the horizon. Directly or indirectly, model applications to active portfolio management depend on the capacity to measure how the portfolio capital requirement changes with changes in portfolio composition.

From a public policy perspective, model-based measurement of economic capital offers a potentially attractive solution to an increasingly urgent regulatory problem. The current regulatory framework for required capital on commercial bank lending is based on the 1988 Basel Accord. Under the Accord, the capital charge on commercial loans is a uniform 8% of loan face value, regardless of the financial strength of the borrower or the quality of collateral.[1] The Accord's rules are risk-sensitive only in that lower charges are specified for certain special classes of lending, e.g., to OECD member governments, to other banks in OECD countries, or for residential mortgages. When the Accord was first introduced, the 8% charge appeared to be "about right on average" for a typical bank portfolio. Over time, however, the failure to distinguish among commercial loans of very different degrees of credit risk created the incentive to move low-risk instruments off balance sheet and retain only relatively high-risk instruments. The financial innovations which arose in response to this incentive have undermined the effectiveness of regulatory capital rules (see, e.g., Jones 2000) and thus led to current efforts towards reform. It is widely recognized that regulatory arbitrage will continue until regulatory capital charges at the instrument level are aligned more closely with underlying risk.

The Basel Committee on Bank Supervision (1999) undertook a detailed study of how banks' internal models might be used for setting regulatory capital. The Committee acknowledged that a carefully specified and calibrated model can deliver a more accurate measure of portfolio credit risk than any rule-based system, but found that the present state of model development could not ensure an acceptable degree of comparability

[1]The so-called 8% rule takes a rather broad definition of "capital." In effect, roughly half this 8% must be in equity capital, as measured on a book-value basis.

across institutions and that data constraints prevent validation of key model parameters and assumptions.[2] It seems unlikely, therefore, that regulators will be prepared in the near- to medium-term to accept the use of internal models for setting regulatory capital. Nonetheless, regulators and industry practitioners appear to be in broad agreement that a revised Accord should permit evolution towards an internal models approach as models and data improve.

At present, it appears virtually certain that a reformed Accord will be a ratings-based "risk-bucketing" system of one form or another. In such a system, banking book assets are grouped into "buckets," which are presumed to be homogeneous. Associated with each bucket is a fixed capital charge per dollar of exposure. At a minimum, one would expect the bucketing system to partition instruments by borrower rating, which would be externally given by rating agencies under some proposals and internally assigned under others; and by one or more proxies for seniority/collateral type, which determines loss severity in the event of default. More complex systems would partition instruments by maturity, country/industry of borrower, and possibly other characteristics. Regardless of the sophistication of the bucketing scheme, capital charges are *portfolio-invariant*, i.e., the capital charge on a given instrument depends only on its own characteristics, and not the characteristics of the portfolio in which it is held. I take portfolio-invariance to be the essential property of risk-bucket capital rules.

A regulatory regime based on risk-bucket assignment of capital charges does offer significant advantages. The current Accord is itself a simple risk-bucketing framework. The reformed Accord could introduce additional bucketing criteria and make better use of information in borrower ratings, yet still be viewed as a natural extension of the current regime. Because the capital charge for a portfolio is simply a weighted sum of the dollars in each bucket, risk-bucketing systems are relatively simple to administer and need not impose burdensome reporting requirements. Validation problems are also limited in scope. Should the use of internal ratings be permitted, the most significant empirical challenge facing supervisors would likely concern the quality of default probability estimates for internal grades.

Though not often recognized in the debate on regulatory reform, in practice many (if not most) large banks apply risk-bucketing for allocation of capital at the transaction level. Even at institutions that have implemented models for portfolio management and portfolio-level capital assessment, there may be reluctance to apply the implied marginal capital requirements to assess hurdle rates for individual transactions. Computational and information systems burdens may be substantial. More important perhaps, line managers are likely to oppose any performance monitoring system in which a loan booked one day at a profitable credit spread becomes unprofitable the next due only to changes in the composition of the bank's overall portfolio. The need for stability in business operations thus favors portfolio-invariant capital charges at the transaction level.

Though risk-bucketing may be a necessary "second-best" solution under current conditions, it is nonetheless desirable that the bucket capital charges be calibrated within a portfolio model. Consistency with a well-specified model would bring greater discipline and accuracy to the calibration process, and would provide

[2] In an industry practitioner response, GARP (1999) acknowledges the obstacles to immediate adoption of an internal models regulatory regime, but argues that the challenges can be met through an evolutionary, piecemeal approach to regulatory certification of model components.

a smoother path of evolution toward an internal-models-based regime. This paper is about the challenges in models-based calibration of risk-bucket capital charges. In particular, it asks what modeling assumptions must be imposed so that marginal contributions to portfolio economic capital are portfolio-invariant.

By design, portfolio models do not, in general, yield portfolio-invariant capital charges. To obtain a distribution of portfolio loss, a model must determine a joint distribution over credit losses at the instrument level. The latest generation of widely-used models gives structure to this problem by assuming that correlations across obligors in credit events arise due to common dependence on a set of *systematic risk factors*. Implicitly or explicitly, these factors represent the sectoral shifts and macroeconomic forces that impinge to a greater or lesser extent on all firms in an economy. A natural property of these models is that the marginal capital required for a loan depends on how it affects diversification, and thus depends on what else is in the portfolio.

If economic capital is defined within the value-at-risk paradigm, then the problem has a simple answer. I show that two conditions are necessary and (with a few regularity conditions) sufficient to guarantee portfolio-invariance: First, the portfolio must be asymptotically fine-grained, in the sense that no single exposure in the portfolio can account for more than an arbitrarily small share of total portfolio exposure. Second, there must be only a single systematic risk factor.

The emphasis in this paper is on generality across portfolios and models. The use of asymptotics to characterize model properties is not new to practitioners, but all previous analyses have been applied to homogeneous portfolios and with the objective of simplifying computation.[3] Banks vary widely in the size and composition of their portfolios and in the details of their credit risk models. For policy purposes, it is essential that our results be sufficiently general to embrace this diversity. Indeed, our results are shown to apply to quite heterogeneous portfolios and across a broad class of credit risk models.

Needless to say, the real world does not give us perfectly fine-grained portfolios. Bank portfolios have finite numbers of obligors and lumpy distributions of exposure sizes. It is clear that capital charges calibrated to the asymptotic case must understate required capital for any given finite portfolio. To assess the magnitude of this bias, I determine the rate of convergence of credit value-at-risk to its asymptotic limit. The results apply generally and appear to be new to the literature. As an application, I propose a simple methodology for assessing a portfolio level add-on charge to compensate for less-than-perfect diversification of idiosyncratic risk. Numerical examples suggest that the method works well, so that departures from asymptotic granularity need not pose a problem in practice for risk-bucket capital rules.

Although it is the standard most commonly applied, value-at-risk is not without shortcomings as a risk-measure for defining economic capital. Because it is based on a single quantile of the loss distribution, VaR provides no information on the magnitude of loss incurred in the event that capital is exhausted. From the perspective of an insurer of deposits (e.g., the FDIC in the US), a more informative summary statistic is *expected excess loss* (“EEL”). Under the EEL paradigm, an institution must hold enough capital so that the expected credit loss in excess of capital is less than or equal to a target loss rate. I consider whether EEL

[3] Large-sample approximations have been applied to homogeneous portfolios under single risk factor versions of the RiskMetrics Group’s CreditMetrics (Finger 1999) and KMV Portfolio Manager (Vasicek 1997) in order to obtain computational shortcuts. Bürgisser, Kurth and Wagner (2000) characterize the asymptotic behavior of a generalized CreditRisk$^+$ model on a sequence of portfolios with n statistically identical copies of a fixed heterogeneous portfolio.

delivers portfolio-invariant capital charges for an asymptotic portfolio in a single-factor setting. Unlike VaR, it does not, and therefore EEL is unsuited as a soundness standard for deriving risk-bucket capital charges.

Section 1 sets out a general framework for the class of risk-factor models in current use under a book-value definition of credit loss. Section 2 presents the key results for VaR for this class of models. In Section 3, these results are shown to apply equally to the case of "multi-state" models in which loss is measured on a market-value basis. The rate of convergence of portfolio VaR to its asymptotic limit is analyzed in Section 4. In Section 5, I show that measures of economic capital based on expected excess loss cannot be made portfolio-invariant, and provide numerical examples of the resulting discrepancies. Concluding remarks focus on the assumption of a single systematic risk factor, which is empirically untenable and yet an unavoidable precondition for portfolio-invariant capital charges. While this assumption ought to be acceptable in the pursuit of achievable and substantive near- to medium-term regulatory reform, it may limit the long-term viability of ratings-based risk-bucket rules for regulatory capital.

1 A general model framework under book-value accounting

Under an *actuarial*, or book-value, definition of loss, credit loss arises only in the event of obligor default. Change in market value due to rating downgrade or upgrade is ignored. This is the simplest framework for our purposes, because we need only be concerned with default risk and with uncertainty in the recovery value of an asset in the event of obligor default.

An essential concept in any risk-factor model is the distinction between *unconditional* and *conditional* event probabilities. An obligor's unconditional default probability, also known as its expected default frequency or *PD*, is the probability of default before some horizon given all information currently observable. The conditional default probability is the PD we would assign the obligor if we also knew what the realized value of the systematic risk factors at the horizon would be. The unconditional PD is the average value of the conditional default probability across all possible realizations of the systematic risk factors. To take an example, consider a simple credit cycle in which the systematic risk factor takes only three values. The "bad state" corresponds to a recession at the risk horizon, the "good state" to an expansion, and the "neutral state" to ordinary times. Say that we currently are in a neutral state, and assign probabilities of $1/4$, $1/2$ and $1/4$ to the three states (respectively) at the risk horizon. Consider an obligor which defaults with probability 2% in the event of a bad state, probability 1% in the neutral state, and probability 0.4% in the event of a good state. The "conditional default probability" is then 0.4%, 1%, or 2%, depending on which horizon state we condition upon. The PD is the probability-weighted average default rate, or 1.1%.

Let X denote the systematic risk factors (possibly multivariate), which are drawn from a known joint distribution. These risk factors may be identified in some models with specific observable quantities, such as macroeconomic variables or industrial sector performance indicators, or may be left abstract. Regardless of their identity, it is assumed that all correlations in credit events are due to common sensitivity to these factors. Conditional on X, the portfolio's remaining credit risk is idiosyncratic to the individual obligors in the portfolio. Let $p_i(x)$ denote the probability of default for obligor i conditional on realization x of X.

This general framework for modeling default is compatible with all of the best-known industry models of portfolio credit risk, including the RiskMetrics Group's CreditMetrics, Credit Suisse Financial Product's CreditRisk$^+$, McKinsey's CreditPortfolioView, and KMV Portfolio Manager. The similarity to CreditRisk$^+$ is easiest to see because that model is written in the language of conditional default probabilities. To obtain CreditRisk$^+$ within our framework, assume that the risk factors $X_1, \ldots, X_K$ are independent gamma-distributed random variables with mean one and variances $\sigma_1^2, \ldots, \sigma_K^2$. Let $\bar{p}_i$ denote the PD of obligor i, and specify $p_i(x)$ as:

$$p_i(x) = \bar{p}_i \left(1 + \sum_{k=1}^{K} w_{ik}(x_k - 1)\right) \tag{1}$$

where w_i is a vector of factor loadings with sum in $[0, 1]$.[4]

CreditMetrics, which is based on a simplified Merton model of default, also can be cast within a conditional probability framework. It is assumed that the vector of risk factors X is jointly distributed $\mathrm{N}(0, \Omega)$. Associated with each obligor is a latent variable R_i which represents the return on the firm's assets. R_i is given by

$$R_i = \xi_i \epsilon_i - X w_i, \tag{2}$$

where the ϵ_i are iid $\mathrm{N}(0, 1)$ white noise (representing obligor-specific risk) and w_i is a vector of factor loadings.[5] Without loss of generality, the weights w_i and ξ_i are scaled so that R_i is mean zero, variance one.[6] A borrower defaults if and only if its asset return falls below a threshold value γ_i.

To obtain the conditional default probability function $p_i(x)$, observe that default occurs if and only if $\epsilon_i \leq (\gamma_i + X w_i)/\xi_i$. Therefore, conditional on $X = x$, default by i is an independent Bernoulli event with probability

$$p_i(x) = \Pr(\epsilon_i \leq (\gamma_i + x w_i)/\xi_i) = \Phi((\gamma_i + x w_i)/\xi_i) \tag{3}$$

where Φ is the standard normal cdf. To calibrate the parameter γ_i, note that the unconditional probability of default is $\Phi(\gamma_i)$, so $\gamma_i = \Phi^{-1}(\bar{p}_i)$, where $\bar{p}_i$ is the PD for obligor i.[7] See Gordy (2000) for a more detailed derivation of these two models and their representation in terms of conditional probabilities.

In some industry models, it is assumed that loss given default ("LGD") is known and non-stochastic. Of the credit VaR models in widespread use, those that do allow for stochastic LGD always take recovery risk to be purely idiosyncratic. In practice, LGD not only may be highly uncertain, but may also be subject to systematic risk. For example, the recovery value of defaulted commercial real estate loans depends on the value of the real estate collateral, which is likely to be lower (higher) when many (few) other real estate

[4] Strictly speaking, this functional form is invalid because it allows conditional probabilities to exceed one. In practice, this problem is negligible for high and moderate quality portfolios and reasonable calibrations of the σ_k^2.

[5] The usual way this is specified has $X w_i$ added, not subtracted. The change in sign here is convenient because it implies that the $p_i(x)$ function will be increasing in x, but does not otherwise change the statistical properties of the model.

[6] Specifically, the weights ξ_i are given by $(1 - w_i' \Omega w_i)^{1/2}$.

[7] By construction, the unconditional distribution of R_i is $\mathrm{N}(0, 1)$, so the probability that $R_i \leq \gamma_i$ is $\Phi(\gamma_i)$.

projects have failed. In recent months, some progress has been made in capturing this effect. Frye (2000) develops an extension of a one-factor CreditMetrics model in which collateral values (and thus recoveries) are correlated with the same systematic risks that drive default rates. Bürgisser et al. (2000) extend the CreditRisk$^+$ model to include a systematic factor for recovery risk that is orthogonal to the systematic factors for default risk.

In order to accommodate systematic and idiosyncratic recovery risk, I take *loss*, rather than merely *default status*, as the primitive outcome variable. Let A_i be the exposure to obligor i; these are taken to be known and non-stochastic.[8] Let the random variable U_i denote loss per dollar exposure. In the event of survival, $U_i = 0$. Otherwise, U_i is the percentage LGD on instrument i, which is assumed to be bounded in the unit interval. The assumption of conditional independence of defaults is extended to conditional independence of the U_i.

For a portfolio of n obligors, define the *portfolio loss ratio* L_n as the ratio of total losses to total portfolio exposure,[9] i.e.,

$$L_n \equiv \frac{\sum_{i=1}^{n} U_i A_i}{\sum_{i=1}^{n} A_i}. \tag{4}$$

For a given $q \in (0,1)$, value-at-risk is defined as the q^{th} percentile of the distribution of loss, and is denoted $\mathrm{VaR}_q[L_n]$. Let $\alpha_q(Y)$ denote the q-th quantile of the distribution of random variable Y, i.e.,

$$\alpha_q(Y) \equiv \inf\{y : \Pr(Y \leq y) \geq q\}. \tag{5}$$

In terms of this more general notation, we have $\mathrm{VaR}_q[L_n] = \alpha_q(L_n)$.

2 Asymptotic loss distribution under book-value accounting

Imagine that the bank selects its portfolio as a large but finite subset of an infinite sequence of lending opportunities. To guarantee that idiosyncratic risk vanishes as more assets are added to the portfolio, the sequence of exposure sizes must neither blow up nor shrink to zero too quickly. I assume that

($\mathcal{A}$-1) the A_i are a sequence of positive constants such that (a) $\sum_{i=1}^{n} A_i \uparrow \infty$ and (b) there exists a $\zeta > 0$ such that $A_n / \sum_{i=1}^{n} A_i = O(n^{-(1/2+\zeta)})$.[10]

The restrictions in ($\mathcal{A}$-1) are sufficient to guarantee that the share of the largest single exposure in total portfolio exposure vanishes to zero. As a practical matter, the restrictions are quite weak and would be

[8] In practice, it need not be so simple. If the instrument is a coupon bond, book-value exposure is simply the face value. Much bank lending, however, is in the form of lines of credit which give the borrower some control over the exposure size. Borrowers do tend to draw down unutilized credit lines as they deteriorate towards default. If we assume that uncertainty in A is idiosyncratic conditional on the state of the obligor and is of bounded variance, then all the conclusions of this paper continue to hold. In this case, we interpret A_i as the *expected* dollar exposure in the event of obligor default.

[9] For simplicity, I assume that the portfolio contains only a single asset for each obligor. Under actuarial treatment of loss, multiple assets of a single obligor may be aggregated into a single asset without affecting the results.

[10] For definition of the order notation $O(\cdot)$ see Definition 2.5 in White (1984).

satisfied by any conceivable real-world large bank portfolio. For example, they are satisfied if all the A_i are bounded from below by a minimum size $A^- > 0$ and from above by a maximum size $A^+ < \infty$.

Our first result is that, under quite general conditions, the conditional distribution of L_n degenerates to its conditional expectation as $n \to \infty$. More formally, we can show that

Proposition 1
If the A_i satisfy ($\mathcal{A}$-1), then, conditional on $X = x$, $L_n - \mathrm{E}[L_n|x] \to 0$, almost surely.

The proof, which relies mainly on a strong law of large numbers, is given in Appendix A. Note that there is no restriction on the relationship between A_i and the distribution of U_i, so there is no problem if, for example, high quality loans tend also to be the largest loans. Also, no restrictions have yet been imposed on the number of systematic factors or their joint distribution.

In intuitive terms, Proposition 1 says that as the exposure share of each asset in the portfolio goes to zero, idiosyncratic risk in portfolio loss is diversified away perfectly. In the limit, the loss ratio converges to a fixed function of the systematic factor X. We refer to this limiting portfolio as "infinitely fine-grained" or as an "asymptotic portfolio." An implication is that, in the limit, we need only know the unconditional distribution of $\mathrm{E}[L_n|X]$ to answer questions about the unconditional distribution of L_n. For example, if we wish to know the variance of the loss ratio, we can look to the variance of $\mathrm{E}[L_n|X]$:

Proposition 2
If the A_i satisfy ($\mathcal{A}$-1), then $\mathrm{V}[L_n] - \mathrm{V}[\mathrm{E}[L_n|X]] \to 0$.

Proof is in Appendix A.

A more important result is, in essence, that for any $q \in (0,1)$, the q^{th} quantile of the unconditional loss distribution approaches the q^{th} quantile of the unconditional distribution of $\mathrm{E}[L_n|X]$ as $n \to \infty$. Our desired result is to have

$$\alpha_q(L_n) - \alpha_q(\mathrm{E}[L_n|X]) \to 0. \tag{6}$$

For technical reasons, however, we are limited to a slightly restricted variant on this result. Let F_n denote the cdf of L_n. We can show:

Proposition 3
If, conditional on $X = x$, $(L_n - \mathrm{E}[L_n|x]) \to 0\, a.s.$ for all x, then for any $\epsilon > 0$

$$F_n(\alpha_q(\mathrm{E}[L_n|X]) + \epsilon) \quad \to \quad [q, 1] \tag{7}$$
$$F_n(\alpha_q(\mathrm{E}[L_n|X]) - \epsilon) \quad \to \quad [0, q]. \tag{8}$$

The proof is in Appendix B. For all practical purposes, this proposition ensures that equation (6) will hold.[11] The literal interpretation of Proposition 3 is that the q^{th} quantile of $\mathrm{E}[L_n|X]$ plus an arbitrarily

[11]The difference has to do with the possibility that the unconditional distributions for the $\{\mathrm{E}[L_n|X]\}$ will permit jump points (or arbitrarily steep slope) at the quantiles $\alpha_q(\mathrm{E}[L_n|X])$ as $n \to \infty$. This possibility is purely a theoretical matter, and would never arise in practical applications.

small "smidgeon" (i.e., ϵ) is guaranteed, in the limit, to cover (or, at least, to come arbitrarily close to covering) q or more of the distribution of loss. Similarly, the q^{th} quantile of $\mathrm{E}[L_n|X]$ less the smidgeon is guaranteed, in the limit, to fail to cover the q^{th} quantile of the distribution of loss (or, at least, to come arbitrarily close to so failing).

The importance of Proposition 3 is that it allows us to substitute the quantiles of $\mathrm{E}[L_n|X]$ (which typically are easy to calculate) for the corresponding quantiles of the loss ratio L_n (which are hard to calculate) as the portfolio becomes large. It should be emphasized that we have obtained this result with very minimal restrictions on the make-up of the portfolio and the nature of credit risk. The assets may be of quite varied PD, expected LGD, and exposure sizes. We have bounded the support of the U_i to the unit interval, but have otherwise not restricted the behavior of the conditional expected loss functions (i.e., the $\mathrm{E}[U_i|x]$).[12] These functions may be discontinuous and non-monotonic, and can vary in form from obligor to obligor. More importantly, we have placed no restrictions on the vector of risk factors X. It may be a vector of any finite length and with any distribution (continuous or discrete).

The quantiles of $\mathrm{E}[L_n|X]$ take on a particularly simple and desirable asymptotic form when we impose two additional restrictions:

($\mathcal{A}$-2) the systematic risk factor X is one-dimensional; and

($\mathcal{A}$-3) the functions $\psi_i(x) \equiv \mathrm{E}[U_i|x]$ are non-decreasing.

These assumptions lead to:

Proposition 4
If ($\mathcal{A}$-2) and ($\mathcal{A}$-3) are satisfied, then $\alpha_q(\mathrm{E}[L_n|X]) = \mathrm{E}[L_n|\alpha_q(X)]$.

Proof: Define

$$\Psi_n(x) = \mathrm{E}[L_n|x] = \frac{\sum_{i=1}^n \psi_i(x)A_i}{\sum_{i=1}^n A_i}. \tag{9}$$

Assumption ($\mathcal{A}$-3) guarantees that $\Psi_n(x)$ is non-decreasing for all x. If $X \leq \alpha_q(X)$, then $\Psi_n(X) \leq \Psi_n(\alpha_q(X))$, so $\Pr(\Psi_n(X) \leq \Psi_n(\alpha_q(X))) \geq \Pr(X \leq \alpha_q(X)) \geq q$. If $\Psi_n(X) < \Psi_n(\alpha_q(X))$, then $X < \alpha_q(X)$, so $\Pr(\Psi_n(X) < \Psi_n(\alpha_q(X))) \leq \Pr(X < \alpha_q(X)) < q$. Therefore,

$$\inf\{y : \Pr(\Psi_n(X) \leq y) \geq q\} = \Psi_n(\alpha_q(X)).$$

QED

Taken together, Propositions 1, 3 and 4 imply a simple and powerful rule for determining capital requirements. For asset i, set capital per dollar book value (inclusive of expected loss) to $c_i \equiv \psi_i(\alpha_q(X)) + \epsilon$, for some arbitrarily small ϵ.[13] Observe that this capital charge depends only on the characteristics of instrument

[12]Technically, the CreditRisk$^+$ model allows U_i to exceed one, because it approximates the Bernoulli distribution of the default event as a Poisson distribution. To accommodate CreditRisk$^+$, we could loosen even this restriction to a requirement that the U_i have bounded variance.

[13]In most practitioner discussions, it is assumed that expected loss is charged against the loan loss reserve and that "capital" refers only to the amount held against unexpected loss. In this paper, "capital" refers to the gross amount set aside.

i and thus this rule is portfolio-invariant. Portfolio losses exceed capital if and only if

$$\sum_{i=1}^{n} U_i A_i > \sum_{i=1}^{n} c_i A_i. \tag{10}$$

Given our rule for c_i and the definition of L_n,

$$\begin{aligned} \Pr\left(\sum_{i=1}^{n} U_i A_i > \sum_{i=1}^{n} c_i A_i\right) &= \Pr\left(L_n > \left(\sum_{i=1}^{n} A_i\right)^{-1} \sum_{i=1}^{n} (\psi_i(\alpha_q(X)) + \epsilon) A_i\right) \\ &= \Pr(L_n > \mathrm{E}[L_n | \alpha_q(X)] + \epsilon) \to [0, 1-q]. \end{aligned}$$

Thus, capital is sufficient, in the limit, so that the probability of portfolio credit losses exceeding portfolio capital is no greater than $1-q$, as desired.

If additional regularity conditions are imposed, the insolvency probability converges to $1-q$ exactly for $\epsilon = 0$. Define the indicator function $\mathcal{I}_i(\delta, B)$ for obligor i, open set B and $\delta > 0$ as equal to one if $\psi_i'(x) > \delta$ for all $x \in B$, and zero otherwise.

($\mathcal{A}$-4) The systematic factor X is continuous, the functions ψ_i are continuous, and there exists $\delta > 0$, $\nu > 0$, $n_0 < \infty$, and an open interval B containing $\alpha_q(X)$ such that

$$\frac{\sum_{i=1}^{n} \mathcal{I}_i(\delta, B) A_i}{\sum_{i=1}^{n} A_i} > \nu, \quad \forall n > n_0.$$

This condition, which in practice would always be satisfied, is sufficient to guarantee that the asymptotic portfolio loss cdf is smooth and increasing in the neighborhood of its q^{th} percentile value, so the technical caveats of Proposition 3 can be circumvented. In Appendix C, I show that

Proposition 5
If assumptions ($\mathcal{A}$-1), ($\mathcal{A}$-2), ($\mathcal{A}$-3) and ($\mathcal{A}$-4) hold, then $\Pr(L_n \leq \mathrm{E}[L_n | \alpha_q(X)]) \to q$.

Therefore, for an infinitely fine-grained portfolio, the proposed portfolio-invariant capital rule provides a solvency probability of exactly q.

Portfolio-invariance depends strongly on the asymptotic assumption and on assumptions ($\mathcal{A}$-2) and ($\mathcal{A}$-3). Portfolios that are not asymptotically fine-grained contain undiversified idiosyncratic risk, which implies that marginal contributions to VaR depend on what else is in the portfolio. Without ($\mathcal{A}$-3), the q^{th} percentile of the distribution of loss need not be associated with the q^{th} percentile of the distribution of X. The quantiles of the loss distribution would then depend in complex ways on the ψ_i functions of the individual obligors. In practice, assumption ($\mathcal{A}$-3) is relatively harmless, because obligors with counter-cyclical credit risk are relatively rare. Indeed, this assumption is imposed by all the well-known latest generation vendor models. Moreover, ($\mathcal{A}$-3) could be relaxed to permit the use of hedging instruments (such as credit derivatives on obligors in the portfolio), so long as the portfolio conditional expected loss function Ψ_n(x) retained the desired monotonicity properties.

Assumption ($\mathcal{A}$-2), however, is much less innocuous from an empirical point-of-view. It is indeed essential to portfolio invariance. If there were, say, two risk factors, and obligors could differ in their sensitivity to each factor, then the realizations (x_1, x_2) associated with a given quantile of the loss distribution would depend on the particular set of obligors in the portfolio. In intuitive terms, the appropriate capital charge for a loan to a heavily-X_1-sensitive borrower would depend on whether the other obligors in the portfolio were predominantly sensitive to X_1 (in which case the loan would add little diversification benefit) or to X_2 (in which case the diversification benefit would be larger). To take a simple example, let X_1 represent the US business cycle and X_2 the European business cycle. Consider the merger of a strictly domestic US, asymptotically fine-grained portfolio with another asymptotically fine-grained bank portfolio. If the second portfolio were also exclusively US, then no diversification benefit would ensue, and required capital for the merged portfolio should be the sum of the capital charges on the two portfolios. However, if the second portfolio contained European obligors, then there would be a diversification benefit (as long as X_1 and X_2 were not perfectly co-monotonic), and the merger should result in reduced total VaR. Therefore, capital charges could not be portfolio-invariant.

Finally, observe that "bucketing" has not appeared, per se, in the derivation. Indeed, the ψ_i functions need not even share a common form across obligors. Sorting obligors into a finite number of statistically homogeneous buckets is helpful for purposes of calibration from data, but is not needed for portfolio-invariant capital charges to be obtained.[14]

3 Asymptotic loss distribution under mark-to-market valuation

Actuarial models are simple to calibrate and understand, and fit naturally with traditional book-value accounting applied to bank loan books. However, much of the credit risk is missed, especially for long-dated highly-rated instruments. Because losses are deemed to arise only in the event of default, no credit loss is recognized when, say, a two-year AA-rated loan downgrades after one year to grade BB. Under a mark-to-market (MTM) notion of loss, credit risk includes the risk of downward (or upward) rating migration, short of default, when the instrument's maturity extends beyond the risk horizon. Even for institutions which report on a book-value basis, it may be desirable to calculate capital charges within a MTM framework in order to capture the additional risk associated with longer instrument maturity.

"Loss" is an ambiguous construct in a mark-to-market setting. I follow one widely-used convention in defining the loss rate U_i on asset i as the difference between expected and realized value at the horizon, discounted by the risk-free rate and divided by current market value.[15] For example, $u_i = 0.2$ represents a 20% loss, and $u_i = -0.05$ represents a 5% gain. Other definitions can be applied without changing the results below. I redefine "exposure" A_i as the current market value.

[14] Multi-state models such as CreditMetrics and CreditPortfolioView typically calibrate PDs to a finite set of rating grades, but the factor loadings w_i may be set at the individual obligor level. In this case, each obligor would comprise its own "bucket." In the KMV model, there is a continuum of "rating grades," so buckets do not arise in any natural way.

[15] Coupon payments, if any, are assumed to be accrued to the horizon at the risk-free rate. Some convention also must be imposed on which intra-horizon cashflows are received on defaulting assets. In practice, how coupons are handled has little effect on the loss distribution, and no qualitative effect on the asymptotics.

Credit risk arises due to uncertainty in U. As before, I assume a vector of systematic risk factors X and that the U_i are conditionally independent. The parameterization and calibration of the $\psi_i(x) \equiv \mathrm{E}[U_i|x]$ functions can draw on existing industry models such as CreditMetrics. Say, for example, that we have a rating system with G non-default grades (grade $G+1$ denoting default), and for each obligor i we have a set of unconditional transition probabilities $\bar{p}_{ig}$ for grade g at the horizon. From these we calculate threshold values γ_{ig} for obligor i's asset return R_i (see equation (2)), such that obligor i defaults if $R_i \leq \gamma_{i,G}$, and transits to "live" grade g if $\gamma_{i,g} < R_i \leq \gamma_{i,g-1}$. The variables $(X, \epsilon_1, \epsilon_2, \ldots, \epsilon_n)$ are iid $\mathrm{N}(0,1)$. Therefore, the conditional transition probabilities are given in CreditMetrics by

$$p_{ig}(x) = \Phi\left((\gamma_{i,g-1} + xw_i)/\sqrt{1-w_i^2}\right) - \Phi\left((\gamma_{i,g} + xw_i)/\sqrt{1-w_i^2}\right), \tag{11}$$

and the unconditional transition probabilities determine the thresholds as $\gamma_{i,g} = \Phi^{-1}(\bar{p}_{i,g-1} + \ldots + \bar{p}_{i,G})$.

Consider a zero-coupon instrument maturing sometime after the horizon. Assume the current value A_i is known, and let $v_{i,g}(x)$ be the value of instrument i at the horizon conditional on the obligor migrating to rating g. In standard implementations of CreditMetrics, pricing at the horizon is done by discounting future contractual cash flows, where the spreads for each grade are taken as fixed and known. In principle, however, we can allow spreads to be non-stochastic functions of X. The conditional expected mark-to-market value at the horizon is

$$\mathrm{MTM}_i(x) = \sum_{g=1}^{G} v_{ig}(x)p_{ig}(x) + \bar{A}_i(1 - \mathrm{E}[\mathrm{LGD}_i|x])p_{i,G+1}(x), \tag{12}$$

where $\bar{A}_i$ is the size of the bank's legal claim on the obligor in the event of a default. Coupons can easily be accommodated in this pricing formula as well with some additional notation. The conditional expected loss functions $\psi_i(x)$ are then given by

$$\psi_i(x) = \frac{\exp(-rT_h)}{A_i}\left(\mathrm{E}[\mathrm{MTM}_i(X)] - \mathrm{MTM}_i(x)\right), \tag{13}$$

where T_h is the time to horizon and r is the risk-free yield for term T_h.

The results of the previous section can be adapted to a mark-to-market setting without difficulty. In contrast to the actuarial case, MTM loss is not bounded from below by zero (e.g., if the obligor's rating improves, there typically will be a gain in value). In principle, it need not be bounded from above either. To guarantee that the convergence properties of Section 2 will continue to hold, I assume that

($\mathcal{A}$-5) all conditional higher moments of loss exist and are bounded; i.e., for each $j \geq 2$, there exists a finite constant μ_j such that $\mathrm{E}[|U_i|^j|x] \leq \mu_j$ for all realizations x and for all instruments i.

For a given portfolio of n assets, L_n, as defined in equation (4), is the discounted portfolio market-valued credit loss at the horizon as a percentage of current market value. I find that all of the Propositions of Section 2 continue to hold, as stated, when ($\mathcal{A}$-5) is imposed as well. Indeed, the proofs in the appendix explicitly accommodate the mark-to-market case. The results is no way depend on the assumptions and conventions of CreditMetrics, which are described above for illustrative purposes.[16] By the same logic as before, the appropriate asymptotic capital charge per dollar current market value for asset i is simply $\psi_i(\alpha_q(X))$.

4 Capital adjustments for undiversified idiosyncratic risk

No portfolio is ever infinitely fine-grained: real-world portfolios have finite numbers of obligors and lumpy distributions of exposure sizes. Large portfolios of consumer loans ought to come close enough to the asymptotic ideal that this issue can safely be ignored, but we ought not to presume the same for even the largest commercial loan portfolios. Unless risk-bucket capital rules are to be abandoned for a full-blown internal models approach, we require a methodology for assessing a capital add-on to cover the residual idiosyncratic risk that remains undiversified in a portfolio. While no simple methodology can be perfect, the same mathematical tools used to analyze the asymptotic behavior of VaR can also provide guidance in constructing a "granularity adjustment" to required capital.

To develop an approximation for the effect of undiversified idiosyncratic risk on portfolio VaR, I analyze how quickly VaR converges to its asymptotic limit as a portfolio grows. This question is most naturally and precisely framed within the context of a homogeneous portfolio, in which each instrument has the same conditional expected loss function $\psi(x)$ and the same exposure size. For this case, I find:

Proposition 6

If the portfolio is statistically homogeneous, the function $\psi(x)$ is continuous, arbitrarily differentiable, and increasing, and the systematic risk factor X is drawn from an arbitrarily differentiable continuous distribution, then VaR converges to its asymptotic limit at rate $1/n$, that is,

$$\alpha_q(L_n) = \psi(\alpha_q(X)) + O(n^{-1}).$$

To prove this proposition, one first shows that the cumulants of the distribution of L_n approach the cumulants of $\mathrm{E}[L_n|X]$ at rate $1/n$. One then applies a generalized Cornish-Fisher expansion to show that the quantiles converge at the same rate. Details are in Appendix D.

Roughly speaking, Proposition 6 says that the difference between the VaR for a given homogeneous portfolio and its asymptotic approximation $\mathrm{E}[L_n|\alpha_q(X)]$ is proportional to $1/n$. This suggests that the accuracy of the asymptotic capital rule can be improved by introducing a "granularity add-on charge" that is inversely proportional to the number of obligors in the portfolio. To calibrate such an add-on charge for a homogeneous portfolio, we need to find a constant of proportionality β such that $\mathrm{E}[L_n|\alpha_q(X)] + \beta/n$ is a good approximation to $\mathrm{VaR}_q[L_n]$.

Of course, Proposition 6 is itself an asymptotic result. When we say that convergence is at rate $1/n$, we are saying that *for large enough* n the gap between VaR and its asymptotic approximation shrinks by half when n is doubled. Short of running the credit VaR model, there is no way to say whether a given n is

[16] In the spirit of KMV Portfolio Manager, for example, one could replace equation (11) with the conditional density function for the default probability at the horizon. The summation in equation (12) would be replaced by an integral, and the v_{ig} would be obtained using risk-neutral valuation. Valuation in the default state in equation (12) also would be modified.

"large enough" for this relationship to hold. To see whether our "$1/n$ rule" works well for realistic values of n and realistic model calibrations, I examine the behavior of VaR in an extended version of CreditRisk$^+$. The virtue of CreditRisk$^+$ for this exercise is that it has an analytic solution. We not only can execute the model for any n very quickly, but also avoid Monte Carlo simulation noise in the results. However, the standard CreditRisk$^+$ model assumes fixed loss given default, and so ignores a potentially important source of volatility.[17] For the special case of a homogeneous portfolio, it is not difficult to augment the model to allow for idiosyncratic recovery risk.

As in the standard CreditRisk$^+$, assume that the systematic risk factor X is gamma-distributed with mean one and variance σ^2. Each obligor has the same default probability $\bar{p}$ and factor loading w. Each facility in the portfolio has identical exposure size, which is normalized to one, and identical expected LGD. The functional form for conditional expected loss function is

$$\psi(x) = \mathrm{E}[\mathrm{LGD}] \cdot \bar{p}(1 + w(x - 1)). \tag{14}$$

To introduce idiosyncratic recovery risk, assume LGD for each obligor is drawn from a gamma distribution with mean λ and variance η^2. This specification is convenient because the sum of m independent and identical gamma random variables is gamma-distributed with mean $m\lambda$ and variance $m\eta^2$. Let G_m denote the gamma cdf with this mean and variance. Let π_m denote the probability that there will be m defaults in the portfolio; these probabilities are calculated in the usual way in CreditRisk$^+$. The probability that $L_n \leq y$ can then be decomposed as

$$\Pr(L_n \leq y) = \sum_{m=0}^{\infty} \pi_m G_m(ny). \tag{15}$$

Long before m approaches n, the π_m become negligibly small, so numerical calculation of equation (15) presents no difficulty. A minor disadvantage of this specification is that it allows LGD to exceed one. However, so long as η is not too large, aggregate losses in the portfolio will be well-behaved, so the problem can be ignored.

Calibration is intended to be qualitatively faithful to available data. When CreditRisk$^+$ is calibrated to rating agency historical performance data, as in Gordy (2000), one finds a negative relationship between $\bar{p}$ and w. By contrast, when a Merton model such as CreditMetrics is calibrated to these data, there is no strong relationship between PD and factor loading. This makes sense, as there is no strong reason to expect that average asset-value correlation should vary systematically across rating grades. To make use of this stylized fact in our calibration, I choose a constant factor loading of $w_{cm} = 0.35$ in CreditMetrics, and calculate a within-grade default correlation for each grade. Shifting back to CreditRisk$^+$, I set a conservative but reasonable value of $\sigma = 2$ for the volatility of X, and then calibrate w for each rating grade so that the within-grade default correlation matches the value from CreditMetrics.[18] The remainder of the calibration

[17] The standard model also implies a discrete loss distribution. As n increases, the "steps" in the loss distribution are re-aligned, which causes local violations of monotonicity in the relationship between n and VaR.

[18] See Gordy (2000) for more details on the choice of σ and on using within-grade default correlations for consistent calibration across the two models.

exercise is straightforward. I choose stylized values for the default probabilities, and assume that LGD has mean 0.4 and standard deviation 0.2. The chosen coverage target is $q = 0.995$ of the loss distribution.

Results are shown in Table 1 for six rating grades. The final column ($n = \infty$) provides the asymptotic capital charge, so the difference between each column and the final column represents the "ideal" granularity add-on. Even for portfolios of only $n = 200$ homogeneous obligors, granularity add-ons are small in the absolute sense (under 60 basis points for the CCC portfolio and under 35 basis points for the investment grade portfolios). However, the add-ons can be large relative to the asymptotic capital charge for investment grade obligors. For a portfolio of 200 homogeneous AA loans, the granularity add-on is roughly double the asymptotic charge.

Figure 1 demonstrates the relationship between the ideal granularity add-on and $1/n$ for each homogeneous portfolio. For the low speculative grades, the predicted linear relationship holds down to $n = 200$.[19] For the high investment grade portfolios, there are small but noticeable departures from the predicted linear relationship when $n < 1000$. The slope of each line (at high values of n) is the appropriate "constant of proportionality" β for the corresponding portfolio. We observe that β is higher for lower quality portfolios, but that values are bounded in a reasonably narrow range (e.g., the appropriate β for a portfolio of B rated facilities is only twice that of a portfolio of A rated facilities). Because departures from linearity are in the concave direction, a granularity adjustment calibrated to the asymptotic slope would slightly overshoot the theoretically optimal add-on for smaller portfolios.

In the case of a non-homogeneous portfolio, determining an appropriate granularity add-on is necessarily more complex and less rigorously founded. However, a simple two-step method appears to work quite well. The first step is to map the actual portfolio to a homogeneous "comparable portfolio" by matching moments of the loss distribution. The second step is to determine the granularity add-on for the comparable portfolio. The same add-on is applied to the capital charge for the actual portfolio.

When the portfolio can be divided into homogeneous buckets, the matching procedure is quite simple and imposes minimal reporting requirements. Now consider a heterogeneous portfolio of n lending facilities divided among B buckets. Within each bucket b, every facility has the same PD $\bar{p}_b$, factor loading w_b, expected LGD λ_b and LGD volatility η_b. Exposure sizes A_i are allowed to vary across facilities in a bucket. To measure the extent to which bucket b exposure is concentrated in a small number of facilities, we require the within-bucket Herfindahl index given by[20]

$$H_b \equiv \frac{\sum_{i \in b} A_i^2}{\left(\sum_{i \in b} A_i\right)^2}.$$

The higher is H_b, the more concentrated is the exposure within the bucket, so the more slowly idiosyncratic risk is diversified away. Finally, let s_b denote the share of total portfolio exposure held in bucket b, i.e.,

$$s_b \equiv \frac{\sum_{i \in b} A_i}{\sum_i A_i}.$$

[19] The slope between each plotted point is constant to six significant digits for both B and CCC portfolios.

[20] The Herfindahl index is a measure of concentration in very widespread use in anti-trust analysis, and should be familiar to many practitioners.

Table 1: Convergence of Quantiles of the Loss Ratio*

	$\bar{p}$	w	$\mathrm{VaR}_q[L_n]$ for values of n: 200	500	1000	1500	2000	3000	5000	∞
AA	0.03	0.933	0.402	0.254	0.198	0.178	0.168	0.158	0.149	0.135
A	0.06	0.850	0.542	0.376	0.315	0.294	0.282	0.271	0.262	0.248
BBB	0.20	0.715	1.039	0.848	0.780	0.756	0.745	0.733	0.723	0.709
BB	1.25	0.526	3.770	3.553	3.475	3.449	3.436	3.423	3.412	3.397
B	6.25	0.369	13.100	12.834	12.745	12.716	12.701	12.686	12.675	12.657
CCC	17.50	0.265	27.937	27.609	27.499	27.463	27.444	27.426	27.411	27.390

*: Default probabilities and VaR expressed in percentage points. Simulations assume $q = 0.995$, $\sigma = 2$, $w_{cm} = 0.35$, $\lambda = 0.4$ and $\eta = 0.2$.

Figure 1: Granularity Add-on as Linear Function of $1/n$

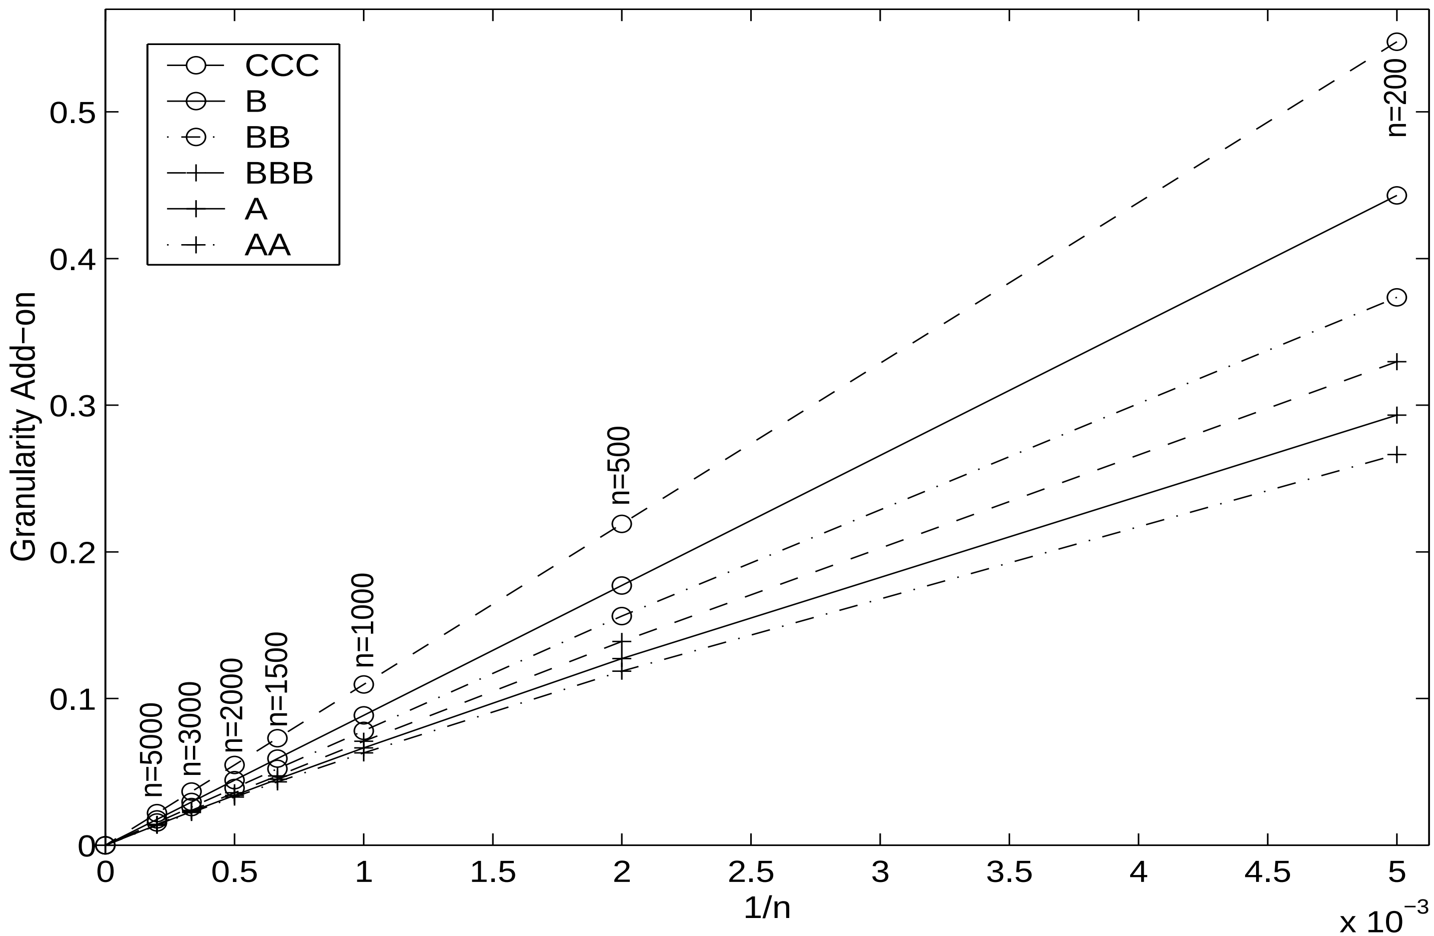

The goal is to construct the comparable portfolio as a portfolio of n^* equal-sized facilities with common PD $\bar{p}^*$, factor loading w^*, and LGD parameters λ^* and η^*. In principle, a wide variety of moment restrictions could be used to do the mapping, but it seems best to choose moments with intuitive interpretation. The first two restrictions equate exposure-weighted expected default rate and expected portfolio loss rate:

$$\bar{p}^* = \sum_{b=1}^{B} \bar{p}_b s_b \qquad \text{and} \qquad \lambda^* \bar{p}^* = \sum_{b=1}^{B} \lambda_b \bar{p}_b s_b. \tag{16}$$

Thus, λ^* is the expected loss rate divided by the expected default rate; i.e.,

$$\lambda^* = \frac{\sum_{b=1}^{B} \lambda_b \bar{p}_b s_b}{\sum_{b=1}^{B} \bar{p}_b s_b}. \tag{17}$$

The remaining moment restrictions equate across the actual and comparable portfolios the contribution to loss variance from different sources of risk. The contribution of systematic risk (i.e., $\mathrm{V}[\mathrm{E}[L|X]]$) takes the simple form

$$\begin{aligned} \mathrm{V}[\mathrm{E}[L_n|X]] &= \sigma^2 \left(\sum_{b=1}^{B} \lambda_b \bar{p}_b w_b s_b \right)^2 \\ \mathrm{V}[\mathrm{E}[L^*|X]] &= \sigma^2 (\lambda^* \bar{p}^* w^*)^2, \end{aligned}$$

which implies

$$w^* = \frac{\sum_{b=1}^{B} \lambda_b \bar{p}_b w_b s_b}{\sum_{b=1}^{B} \lambda_b \bar{p}_b s_b}. \tag{18}$$

Note that w^* is simply an expected-loss-weighted average of the w_b.

The contribution of idiosyncratic risk to loss variance (i.e., $\mathrm{E}[\mathrm{V}[L_n|X]]$) works out to

$$\begin{aligned} \mathrm{E}[\mathrm{V}[L_n|X]] &= \sum_{b=1}^{B} \left(\lambda_b^2 (\bar{p}_b(1-\bar{p}_b) - (\bar{p}_b w_b \sigma)^2) + \bar{p}_b \eta_b^2 \right) H_b s_b^2 \\ \mathrm{E}[\mathrm{V}[L^*|X]] &= \frac{1}{n^*} \left(\lambda^{*2} (\bar{p}^*(1-\bar{p}^*) - (\bar{p}^* w^* \sigma)^2) + \bar{p}^* \eta^{*2} \right). \end{aligned}$$

Terms containing $\lambda^2(\bar{p}(1-\bar{p}) - (\bar{p}w\sigma)^2)$ represent the contribution of idiosyncratic default risk, and terms containing $\bar{p}\eta^2$ represent the contribution of idiosyncratic recovery risk. By matching these two contributions separately, I get the final two restrictions needed for identification. The number of exposures in the comparable portfolio works out to

$$n^* = \left(\sum_{b=1}^{B} \Lambda_b H_b s_b^2 \right)^{-1} \tag{19}$$

where

$$\Lambda_b \equiv \frac{\lambda_b^2(\bar{p}_b(1-\bar{p}_b) - (\bar{p}_b w_b \sigma)^2)}{\lambda^{*2}(\bar{p}^*(1-\bar{p}^*) - (\bar{p}^* w^* \sigma)^2)}.$$

The form of equation (19) allows n^* to be interpreted as an inverse measure of weighted exposure concentration. Finally, the variance of LGD for the comparable portfolio is given by

$$\eta^{*2} = \frac{n^*}{\bar{p}^*} \sum_{b=1}^{B} \eta_b^2 \bar{p}_b H_b s_b^2. \qquad (20)$$

The remaining step is to determine the constant of proportionality β^* for the comparable homogeneous portfolio. In principle, this constant ought to be estimated directly for each parameter set $(\bar{p}^*, w^*, \lambda^*, \eta^*)$. The necessary calculations could be automated in a spreadsheet, and should require very little time to execute. In practice, however, the need for operational transparency dictates that we calibrate a simple rule for β^* as a function of $(\bar{p}^*, w^*, \lambda^*, \eta^*)$. Assuming that bucket factor loadings and LGD volatilities are themselves calibrated as functions of bucket PD and expected LGD, respectively, the rule can perhaps be expressed as a lookup table mapping $(\bar{p}^*, \lambda^*)$ to β^*.

Matching lower-order moments gives no guarantee that the loss distribution for the comparable portfolio will display higher-order moments very close to those of the original heterogeneous portfolio. Tail quantiles of the loss distribution are sensitive to higher-order moments, so the performance of the methodology needs to be confirmed on a range of empirically plausible portfolios. As an example, I construct a portfolio of 1000 obligors divided equally across four buckets. The buckets represent high investment grade, low investment grade, high speculative grade and moderate to low speculative grade. Factor loadings are calibrated as in Table 1. Expected LGDs for the buckets are set to 0.2, 0.3, 0.4, and 0.5, respectively, and the LGD volatility is set (rather generously) to $\eta_b = \sqrt{\lambda_b(1-\lambda_b)}$. Table 2 displays the bucket-level parameters.

Table 2: Bucket-level Parameters of Stylized Portfolio*

	$\bar{p}$	w	λ	η
1	0.05	0.871	0.2	0.400
2	0.50	0.618	0.3	0.458
3	1.00	0.548	0.4	0.490
4	5.00	0.391	0.5	0.500

*: Default probabilities in percentage points.

Exposure size for facility i is set to i^4; i.e., A_i is \$1 for the first exposure, \$16 for the second, \$81 for the third, and so on. The exposures are assigned to buckets by turn. The first exposure is assigned to Bucket 4, the second to Bucket 3, the third to Bucket 2, the fourth to Bucket 1, the fifth to Bucket 4, and so on. Looking at the portfolio as a whole, I find that the largest 10% of exposures account for roughly 40% of total exposure, which matches the empirical rule of thumb reported by Carey (forthcoming) for concentration of outstandings. Also, portfolio exposure is roughly split between investment and speculative grades, which appears to be typical of a commercial loan portfolio at a large bank.[21]

I first obtain parameters for the comparable homogeneous portfolio. The comparable portfolio has $n^* = 372.0$ obligors, which is under 40% of the obligor count of the original portfolio.[22] Each obligor has PD of 1.63% and factor loading $w^* = 0.4245$. Loss given default has expected value 0.467 and volatility 0.502. By construction, the comparable portfolio matches the original portfolio in its expected loss rate of 0.760%. For each portfolio, the standard deviation of the loss rate is 0.787%.

I next obtain the granularity add-on parameter c for the comparable portfolio. For the purpose of this illustration, I use equation (15) to calculate the loss distribution for homogeneous portfolios of $n = 3000$ and $n = 4000$ facilities with parameters $(\bar{p}^*, w^*, \lambda^*, \eta^*)$, then estimate the slope parameter β^* for coverage target q as

$$\beta_q^* = \frac{\mathrm{VaR}_q[L_{3000}] - \mathrm{VaR}_q[L_{4000}]}{1/3000 - 1/4000}. \tag{21}$$

Finally, I approximate VaR_q for the original portfolio as its asymptotic VaR plus β_q^*/n^*.

Results are shown in Table 3 for three tail values of q. Row (*i*) presents estimates of VaR obtained by direct simulation of the original portfolio. Row (*ii*) presents the asymptotic VaR for the original portfolio given by $\mathrm{E}[L_n|\alpha_q(X)]$. Row (*iii*) shows VaR for the comparable portfolio obtained using equation (15). The granularity add-on β_q^*/n^* is shown in row (*iv*). Row (*v*) sums the asymptotic VaR and granularity add-on to get our approximation. Tracking error between rows (*v*) and (*i*) is shown in the final row.

The procedure works well for all values of q. The error due to linear approximation of the "$1/n$ rule" is minimal. At $q = 99.5\%$, the linear approximation matches the directly calculated VaR for the comparable portfolio to four decimal places. The error due to using the comparable portfolio in place of the original portfolio is also small. At $q = 99.5\%$, our approximated VaR overshoots its target by 4 basis points.

Table 3: Direct and Approximated Estimates of VaR*

	q:	99.0	99.5	99.9
(*i*)	"True" VaR	3.952	4.713	6.641
(*ii*)	Asymptotic VaR	3.579	4.312	6.086
(*iii*)	VaR for comparable portfolio	3.944	4.752	6.708
(*iv*)	Granularity add-on	0.356	0.440	0.622
(*v*)	Approximated VaR	3.935	4.752	6.708
(*vi*)	Tracking error	−0.017	+0.040	+0.066

*: All quantities expressed in percentage points. "True" VaR estimated by simulation with 200,000 Monte Carlo trials.

For the special case of the CreditRisk$^+$ model, the portfolio data needed for the moment-matching map-

[21]In a sample of large bank commercial loan portfolios, Treacy and Carey (1998, Chart 3) show that roughly half of aggregate internally rated outstandings are investment grade.

[22]Note that the procedure for calculating the granularity add-on does not require n^* to be an integer.

ping method should pose minimal additional reporting burden for regulated institutions. Default probability, expected LGD and total bucket exposure would need to be reported by the bank to calculate the asymptotic capital charge. The factor loadings and (presumably) LGD volatilities are to be chosen by the regulator. The only new required inputs, the within-bucket Herfindahl indices, are easily calculated from the individual exposure sizes.

It should be emphasized that the theoretical underpinnings for the granularity adjustment apply equally to mark-to-market models. The simple linear formulae for parameters of the comparable homogeneous portfolio depend on the linear functional forms assumed in CreditRisk$^+$. Specifications based on more complex models, e.g., KMV Portfolio Manager or CreditMetrics, imply more complex mapping formulae whose inputs need not be reducable to bucket-level summary statistics (e.g., Herfindahl indices). However, it seems reasonable to conjecture that one can achieve tolerable accuracy using crude rules based on the CreditRisk$^+$ formulae. What is most important is that there be a reasonably accurate measure for the "effective" obligor count (i.e., n^*) in a heterogeneous portfolio. Most bank portfolios are heavy-tailed in exposure size distribution, and thus may have an effective n^* that is an order of magnitude smaller than the raw obligor count in the portfolio.

5 Asymptotic properties of EEL capital charges

Industry application of credit risk modeling to capital allocation appears invariably to equate soundness with a coverage target for value-at-risk. The VaR soundness standard does have shortcomings, at least in principle. Because it ties capital to a single percentile of the distribution of loss, it can lead to "cliff effects." To take an extreme example, say we construct a portfolio in which losses are all or nothing. That is, the loss rate equals zero with probability p, and equals one with probability $1 - p$. If the VaR target quantile q is less than or equal to p, the VaR capital charge is zero, but if $q = p + \epsilon$ for any positive ϵ, the capital charge is 100%. Banks potentially could securitize portfolios into structured tranches in order to arbitrage this effect. A closely related problem is that VaR can exhibit a counterintuitive non-monotonic relationship with the variance and higher moments of the loss distribution. If factor loadings are pushed high enough, then the probability mass of the loss distribution can be pushed so far into the tail (beyond the q^{th} percentile) that further increasing factor loadings actually decreases VaR. Thus, VaR can decline as the probability of a cataclysmic loss increases.

As an alternative to VaR, a soundness standard can be based on *expected excess loss* ("EEL"). Under the EEL paradigm, an institution holds capital (plus reserves) so that the expected credit loss in excess of capital is less than or equal to a target loss rate θ. That is, if θ is the targeted loss rate and L is credit loss as a percentage of total exposure, then the required capital (plus reserves) is c per dollar of total exposure, where c solves

$$\theta = \text{EEL}[L, c] \equiv \text{E}[(L - c)^+], \tag{22}$$

and where Y^+ denotes $\max(Y, 0)$.

EEL offers some important advantages as a soundness standard. Cliff effects are eliminated. If the distribution of L is continuous at c, then the EEL operator is continuous with respect to c, so small changes in the θ standard can be accommodated by small changes in c.[23] In all the models in widespread use, either the distribution of L is continuous or it becomes continuous in the asymptotic limit. McAllister and Mingo (1996) propose EEL-based soundness standards for securitizations to prevent banks from exploiting the VaR cliff effect. EEL also avoids the "back-bending" problem. As long as c is greater than expected loss, any mean-preserving spread in the distribution of loss increases EEL, so c must increase in order to maintain EEL at its target level.

Unfortunately, risk-bucket capital charges cannot be derived under an EEL paradigm, because EEL capital charges are never portfolio-invariant. Some intuition for this problem can be gained by writing the asymptotic EEL for homogeneous portfolios in terms of the distribution of the systematic risk factor. Assume we have loans of two types, denoted "a" and "b". Let $\psi_a(x)$ denote the expected loss for bucket a loans conditional on $X = x$. Under regulatory conditions ($\mathcal{A}$-1) and ($\mathcal{A}$-5), $L_a - \psi_a(X) \to 0$, almost surely, so

$$\mathrm{EEL}[L_a, c_a] = \mathrm{E}[(L_a - c_a)^+] \to \mathrm{E}[(\psi_a(X) - c_a)^+]. \tag{23}$$

The asymptotic EEL capital charge c_a is chosen to set this quantity to the desired target θ. Similar analysis for bucket b gives c_b.

Now say we have a mixed portfolio containing equal numbers of loans from a and b. For simplicity, the exposures are equal-sized. Asymptotic EEL for the mixed portfolio is given by $\mathrm{E}[(\psi_m(X) - c_m)^+]$. By construction of the mixed portfolio, we have $\psi_m(X) = (\psi_a(X) + \psi_b(X))/2$. If asymptotic EEL were portfolio-invariant, then a capital charge of $c_m = (c_a + c_b)/2$ would yield an EEL of θ for the mixed portfolio. To test whether this holds, I evaluate

$$\mathrm{E}[(\psi_m(X) - c_m)^+] = \mathrm{E}[((\psi_a(X) - c_a)/2 + (\psi_b(X) - c_b)/2)^+]. \tag{24}$$

We now require the following lemma:

Lemma 1
If Y_1 and Y_2 are integrable random variables on a probability space $(\Omega, \mathcal{F}, P)$, then

$$\mathrm{E}[(Y_1 + Y_2)^+] \leq \mathrm{E}[Y_1^+] + \mathrm{E}[Y_2^+]. \tag{25}$$

If $P(\{\omega : (Y_1(\omega) < 0 < Y_2(\omega)) \vee (Y_2(\omega) < 0 < Y_1(\omega))\}) > 0$, then the inequality in equation (25) is strict.

Proof is given in Appendix E.

[23]More formally, if $C(\theta)$ is a function mapping θ to the solution for c in equation (22), then by taking total derivatives we get $C'(\theta) = -1/(1 - F_L(C(\theta)))$, where F_L is the cdf of L.

The conditions of Lemma 1 apply to $Y_i \equiv (\psi_i(X) - c_i)/2$, which gives us

$$\mathrm{E}[(\psi_m(X) - c_m)^+] \leq \mathrm{E}[((\psi_a(X) - c_a)/2)^+] + \mathrm{E}[((\psi_b(X) - c_b)/2)^+] = \theta. \tag{26}$$

In general, when EEL is fixed across buckets, the threshold realization of X at which homogeneous portfolio a hits insolvency does not equal the corresponding threshold for homogeneous portfolio b, so for some interval of x values we will have either $\psi_a(x) - c_a < 0 < \psi_b(x) - c_b$ or $\psi_a(x) - c_a > 0 > \psi_b(x) - c_b$.[24] Therefore, the inequality in equation (26) will in most situations be strict, which implies that c_m is too strict a capital requirement for the asymptotic mixed portfolio.

To provide a rough idea of how much we overshoot required capital in a mixed portfolio, I apply EEL to an asymptotic, single systematic factor version of CreditRisk$^+$. In Appendix F, I show that asymptotic EEL takes on a relatively simple form in this model. Table 4 presents EEL- and VaR-based capital requirements for homogeneous asymptotic portfolios of different credit ratings. Parameters for each rating grade and the volatility of X are taken from Table 1. The "EEL" and "VaR" columns in Table 4 report required capital charges (gross of reserves) for an EEL target of $\theta = 0.00002$ (i.e., 0.2 basis points) and a VaR target of $q = 99.5\%$, respectively. The value of θ was chosen to equate capital requirements under the two standards for an obligor at the border of investment and speculative grades. In this example, the EEL standard produces lower (higher) capital requirements than VaR for the higher (lower) grades. Without further analysis, it is unclear whether this result is model-dependent.

Table 4: Asymptotic EEL and VaR Capital Charges*

	EEL	VaR
AA	0.050	0.135
A	0.131	0.248
BBB	0.571	0.709
BB	4.135	3.397
B	19.352	12.657
CCC	45.550	27.390

*: Capital in percentage points.

I next form mixed portfolios. In each case, I assume an asymptotic portfolio of equal-sized loans, half of which are in one bucket and half in another bucket. It is straightforward to show that the conditional expected loss rate for a mixed portfolio is

$$\psi_m(x) = \frac{1}{2}\psi_a(x) + \frac{1}{2}\psi_b(x) = \lambda \bar{p}_m(1 - w_m + w_m x) \tag{27}$$

where $\bar{p}_m = (\bar{p}_a + \bar{p}_b)/2$ and $w_m = (\bar{p}_a w_a + \bar{p}_b w_b)/(2\bar{p}_m)$. The $\psi_m(x)$ take on the same form as for homogeneous portfolios, so the tools of Appendix F apply without modification. Results for four different mixed portfolios are presented in Table 5. The third column shows the capital charge required for the

[24] These threshold values for X are given by $x_i = \psi_i^{-1}(c_i)$.

mixed portfolio to hit the EEL target of θ. The fourth column shows the average of the capital charges for homogeneous portfolios of the two constituent buckets. The final column shows the "tracking error" as a percentage of the third column. As one would expect, the average of the homogeneous capital charges overshoots the correct mixed-portfolio capital charge by a relatively small (though non-negligible) amount when the two buckets are adjacent. For a mix of grades AA and A, we overshoot by under 3%. For a mix of BBB and BB, we overshoot by 6.5%. If distant buckets are mixed, the overshoot is much larger (over 16% for the two examples in the table).

Table 5: Asymptotic EEL for Mixed Portfolios

Bucket a	Bucket b	c_m	$(c_a + c_b)/2$	Error
AA	A	0.088	0.090	+2.7%
A	B	8.378	9.741	+16.3%
BBB	BB	2.210	2.353	+6.5%
AA	CCC	19.658	22.800	+16.0%

*: Capital charges expressed in percentage points.

Discussion

This paper shows how risk-factor models of credit value-at-risk can be used to justify and calibrate a ratings-based risk-bucket system for assigning capital charges for credit risk at the instrument level. Risk-bucket systems, by definition, permit capital charges to depend only on the characteristics of the instrument and its obligor, and not the characteristics of the remainder of the portfolio. Risk-factor models deliver this property, which I call *portfolio invariance,* only if two conditions are satisfied. First, the portfolio must be asymptotically fine-grained, in order that all idiosyncratic risk be diversified away. Second, there can be only a single systematic risk factor.

Violation of the first condition, which occurs for every finite portfolio, does not pose a serious obstacle in practice. Analysis of rates of convergence of VaR to its asymptotic limit leads to a theoretically sound and practical method of approximating a portfolio-level adjustment for undiversified idiosyncratic risk.

The second condition presents a greater dilemma. The single risk factor assumption, in effect, imposes a single monolithic business cycle on all obligors. A revised Basel Accord must apply to the largest international banks, so the single risk factor should in principle represent the global business cycle. By assumption, all other credit risk is strictly idiosyncratic to the obligor. In reality, the global business cycle is a composite of a multiplicity of cycles tied to geography (e.g., political shifts, natural disasters) or to prices of production inputs (e.g., oil, commodities). A single factor model cannot capture any clustering of firm defaults due to common sensitivity to these smaller-scale components of the global business cycle. Holding fixed the state of the global economy, a local recession in, for example, Spain is permitted to contribute nothing to the default rate of Spanish obligors. If there are indeed pockets of risk, then calibrating a single factor model to a broadly diversified international credit index may significantly understate the capital needed to support a regional or specialized lender.

Would empirical violation of the single factor assumption necessarily render a risk-bucket capital rule unreliable and ineffective? The answer depends on the scope of application and the sophistication of debt markets. Regulators will need to use caution and judgement in applying risk-bucket capital charges to institutions that are less broadly diversified. One should note that the current Basel Accord, which is itself a risk-bucket system, is applied to an enormous range of institutions, so it seems unlikely that a reformed Accord would bring about any greater harm.

More generally, the ability of banks to subvert risk-bucket capital rules by exploiting the inadequacy of the single factor assumption depends on the capacity of debt markets to recognize and price different risk-factors. At present, this level of sophistication appears to be lacking. Partly because markets do not yet provide precise information on correlations of credit events across obligors, many (perhaps most) of the institutions that actively use credit VaR models effectively impose the single-factor assumption.[25] In the near- to medium-term, therefore, the implausibility of the single factor assumption need not present an obstacle to the implementation of reformed ratings-based risk-bucket capital rules. In the long run, however, the need to relax this assumption may impel adoption of a more sophisticated internal-models regulatory regime.

Appendix

A Proof of Propositions 1 and 2

The proof of Proposition 1 requires a version of the the strong law of large numbers for a sequence $\{Y_n\}$ of random variables and a sequence $\{a_n\}$ of positive constants:

Lemma 2
If $a_n \uparrow \infty$ and $\sum_{n=1}^{\infty}(\mathrm{V}[Y_n]/a_n^2) < \infty$, then

$$\left(\sum_{i=1}^{n} Y_n - \mathrm{E}[\sum_{i=1}^{n} Y_n]\right)/a_n \to 0\, a.s..$$

Proof is given by Petrov, ed (1995), Theorem 6.7. We also make use of the following lemma:

Lemma 3
If $\{b_n\}$ is a sequence of positive real numbers such that $\{b_n\}$ is $O(n^{-\rho})$ for some $\rho > 1$, then $\sum_{n=1}^{\infty} b_n < \infty$.

This lemma is a corollary of Theorem 3.5.2 in Knopp (1956) and the convergence of the harmonic series $1/n^{\rho}$ for $\rho > 1$ (see Knopp 1956, Example 3.1.2.3).

[25] Users of KMV Portfolio Manager and CreditMetrics often impose a uniform asset-value correlation across obligors. Users of CreditRisk$^+$ typically assume a single factor and a factor loading of $w = 1$ for all obligors. In both these examples, the user is implicitly imposing both a single systematic factor and a uniform value for the factor loading.

We now prove Proposition 1. Let $Y_n \equiv U_n A_n$ and $a_n \equiv \sum_{i=1}^{n} A_i$. conditional on $X = x$, we have

$$\sum_{n=1}^{\infty} (\mathrm{V}[Y_n]/a_n^2) = \sum_{n=1}^{\infty} \left(A_n / \sum_{i=1}^{n} A_i \right)^2 \mathrm{V}[U_n|x]$$

Under the actuarial definition of loss, U_n is bounded in $[0,1]$, so we must have $\mathrm{V}[U_n|x] < 1$ for any $X = x$. For this proposition to hold under the mark-to-market paradigm as well, we require assumption ($\mathcal{A}$-5), which implies

$$\mathrm{V}[U_n|x] \leq \mathrm{E}[|U_n|^2|x] \leq \mu_2.$$

Therefore, under either definition of loss, there exists a finite constant μ_2 such that

$$\sum_{n=1}^{\infty} (\mathrm{V}[Y_n]/a_n^2) \leq \mu_2 \sum_{n=1}^{\infty} \left(A_n / \sum_{i=1}^{n} A_i \right)^2 .$$

By part (b) of assumption ($\mathcal{A}$-1), the sequence $\{A_n / \sum_{i=1}^{n} A_i\}$ is $O(n^{-(1/2+\zeta)})$ for some $\zeta > 0$, so the sequence $\left\{ (A_n / \sum_{i=1}^{n} A_i)^2 \right\}$ is $O(n^{-(1+2\zeta)})$. By Lemma 3, the series sum must be finite. By part (a) of assumption ($\mathcal{A}$-1), we have $a_n \uparrow \infty$. The conditions of Lemma 2 are thus satisfied. The loss ratio L_n is equal to $\sum_{i=1}^{n} Y_i / a_n$, so Proposition 1 is proved. **QED**

Proposition 2 follows similar logic. We require the following lemma:

Lemma 4

Let $\{b_n\}$ *and* $\{d_n\}$ *be sequences of real numbers such that* $a_n \equiv \sum_{i=1}^{n} b_i \uparrow \infty$ *and* $d_n \rightarrow 0$. *Then* $(1/a_n) \sum_{i=1}^{n} b_i d_i \rightarrow 0$.

This result is a special case of Petrov, ed (1995), Lemma 6.10. If we let $b_n = A_n$ and $d_n = A_n / \sum_{i=1}^{n} A_i$, then assumption ($\mathcal{A}$-1) guarantees that $a_n \uparrow \infty$ and $d_n \rightarrow 0$, so apply Lemma 4 to get

$$\frac{1}{\sum_{i=1}^{n} A_i} \sum_{i=1}^{n} \frac{A_i^2}{\sum_{j=1}^{i} A_j} \rightarrow 0. \tag{28}$$

The standard rule for conditional variance gives

$$\mathrm{V}[L_n] - \mathrm{V}[\mathrm{E}[L_n|X]] = \mathrm{E}[\mathrm{V}[L_n|X]] = \mathrm{E}\left[\frac{\sum_{i=1}^{n} A_i^2 \mathrm{V}[U_i|X]}{(\sum_{i=1}^{n} A_i)^2} \right].$$

$\mathrm{E}[\mathrm{V}[U_i|X]]$ must be less than one under the actuarial paradigm and is bounded from above by assumption ($\mathcal{A}$-5) under the mark-to-market paradigm. Therefore, there exists a finite constant μ_2 such that

$$\mathrm{E}[\mathrm{V}[L_n|X]] \leq \mu_2 \frac{\sum_{i=1}^{n} A_i^2}{(\sum_{i=1}^{n} A_i)^2} = \mu_2 \frac{1}{\sum_{i=1}^{n} A_i} \sum_{i=1}^{n} \frac{A_i^2}{\sum_{j=1}^{n} A_j} \leq \mu_2 \frac{1}{\sum_{i=1}^{n} A_i} \sum_{i=1}^{n} \frac{A_i^2}{\sum_{j=1}^{i} A_j} \rightarrow 0.$$

As $\mathrm{E}[\mathrm{V}[L_n|X]]$ must be non-negative and is bounded from above by quantity converging to zero, it too must converge to zero. **QED**

B Proof of Proposition 3

Almost sure convergence implies convergence in probability (see White 1984, Theorem 2.24), so for all x and $\epsilon > 0$,

$$\Pr(|L_n - \mathrm{E}[L_n|x]| < \epsilon|x) \to 1. \tag{29}$$

If F_n is the cdf of L_n, then equation (29) implies

$$F_n(\mathrm{E}[L_n|x] + \epsilon|x) - F_n(\mathrm{E}[L_n|x] - \epsilon|x) \to 1.$$

Because F_n is bounded in $[0, 1]$, we must have $F_n(\mathrm{E}[L_n|x] + \epsilon|x) \to 1$ and $F_n(\mathrm{E}[L_n|x] - \epsilon|x) \to 0$.

Let S_n^+ denote the set of realizations x of X such that $\mathrm{E}[L_n|x]$ is less than or equal to its q^{th} quantile value, i.e.,

$$S_n^+ \equiv \{x : \mathrm{E}[L_n|x] \leq \alpha_q(\mathrm{E}[L_n|X])\}.$$

By construction, $\Pr(x \in S_n^+) \geq q$.

By the usual rules for conditional probability, we have

$$\begin{aligned} F_n(\alpha_q(\mathrm{E}[L_n|X]) + \epsilon) &= F_n(\alpha_q(\mathrm{E}[L_n|X]) + \epsilon|X \in S_n^+) \Pr(X \in S_n^+) \\ &\quad + F_n(\alpha_q(\mathrm{E}[L_n|X]) + \epsilon|X \notin S_n^+) \Pr(X \notin S_n^+) \\ &\geq F_n(\alpha_q(\mathrm{E}[L_n|X]) + \epsilon|X \in S_n^+) \Pr(X \in S_n^+) \\ &\geq F_n(\alpha_q(\mathrm{E}[L_n|X]) + \epsilon|X \in S_n^+)\, q \end{aligned} \tag{30}$$

For all $x \in S_n^+$, we have

$$F_n(\alpha_q(\mathrm{E}[L_n|X]) + \epsilon|x) \geq F_n(\mathrm{E}[L_n|x] + \epsilon|x) \to 1$$

so the dominated convergence theorem implies that[26]

$$F_n(\alpha_q(\mathrm{E}[L_n|X]) + \epsilon|X \in S_n^+) \to 1,$$

so from equation (30) we have

$$F_n(\alpha_q(\mathrm{E}[L_n|X]) + \epsilon) \to [q, 1]$$

[26] See Theorem 16.4 in Billingsley (1995).

as required.

The other half of the proof follows similarly. Define S_n^- as

$$S_n^- \equiv \{x : \mathrm{E}[L_n|x] \geq \alpha_q(\mathrm{E}[L_n|X])\}$$

so that $\Pr(x \in S_n^-) \geq 1-q$. Then

$$\begin{aligned} F_n(\alpha_q(\mathrm{E}[L_n|X]) - \epsilon) &= F_n(\alpha_q(\mathrm{E}[L_n|X]) - \epsilon | X \notin S_n^-) \Pr(X \notin S_n^-) \\ &\quad + F_n(\alpha_q(\mathrm{E}[L_n|X]) - \epsilon | X \in S_n^-) \Pr(X \in S_n^-) \\ &\leq q + F_n(\alpha_q(\mathrm{E}[L_n|X]) - \epsilon | X \in S_n^-) \Pr(X \in S_n^-). \end{aligned} \tag{31}$$

For all $x \in S_n^-$, we have

$$F_n(\alpha_q(\mathrm{E}[L_n|X]) - \epsilon | x) \leq F_n(\mathrm{E}[L_n|x] - \epsilon | x) \to 0$$

so the dominated convergence theorem implies that

$$F_n(\alpha_q(\mathrm{E}[L_n|X]) - \epsilon | X \in S_n^-) \to 0,$$

so from equation (31) we have

$$F_n(\alpha_q(\mathrm{E}[L_n|X]) - \epsilon) \to [0, q]$$

as required. **QED**

C Proof of Proposition 5

The proof of Proposition 5 requires the following lemma:

Lemma 5
Let Y_1 and Y_2 be random variables with cdfs F_1 and F_2, respectively. For all y and all $\epsilon > 0$,

$$|F_1(y) - F_2(y)| \leq \Pr(|Y_1 - Y_2| > \epsilon) + \max\{F_2(y+\epsilon) - F_2(y), F_2(y) - F_2(y-\epsilon)\}.$$

Proof: For any y and any $\epsilon > 0$,

$$\begin{aligned} \{Y_1 : Y_1 \leq y\} &\subset \{Y_2 : Y_2 \leq y + \epsilon\} \cup \{(Y_1, Y_2) : |Y_1 - Y_2| > \epsilon\} \\ \{Y_2 : Y_2 \leq y - \epsilon\} &\subset \{Y_1 : Y_1 \leq y\} \cup \{(Y_1, Y_2) : |Y_1 - Y_2| > \epsilon\}, \end{aligned}$$

so we must have

$$\begin{aligned}
\Pr(Y_1 \leq y) &\leq \Pr(Y_2 \leq y+\epsilon) + \Pr(|Y_1 - Y_2| > \epsilon) \\
\Pr(Y_1 \leq y) &\geq \Pr(Y_2 \leq y-\epsilon) - \Pr(|Y_1 - Y_2| > \epsilon).
\end{aligned}$$

Substract $F_2(y)$ from both sides of each equation, and rearrange to get

$$\begin{aligned}
F_1(y) - F_2(y) &\leq F_2(y+\epsilon) - F_2(y) + \Pr(|Y_1 - Y_2| > \epsilon) \\
F_2(y) - F_1(y) &\leq F_2(y) - F_2(y-\epsilon) + \Pr(|Y_1 - Y_2| > \epsilon).
\end{aligned}$$

Combine these inequalities to complete the proof. **QED**

To apply Lemma 5, let $Y_1 = L_n$ and $Y_2 = \mathrm{E}[L_n|X]$, and let F_n^* denote the cdf of $\mathrm{E}[L_n|X]$. We then have

$$\begin{aligned}
&|F_n(\mathrm{E}[L_n|\alpha_q(X)]) - F_n^*(\mathrm{E}[L_n|\alpha_q(X)])| \leq \Pr(|L_n - \mathrm{E}[L_n|X]| > \epsilon) \\
&+ \max\{F_n^*(\mathrm{E}[L_n|\alpha_q(X)] + \epsilon) - F_n^*(\mathrm{E}[L_n|\alpha_q(X)]), F_n^*(\mathrm{E}[L_n|\alpha_q(X)]) - F_n^*(\mathrm{E}[L_n|\alpha_q(X)] - \epsilon)\} \quad (32)
\end{aligned}$$

for any $\epsilon > 0$.

The regularity conditions in ($\mathcal{A}$-4) are sufficient to guarantee that the cdf F_n^* is differentiable in the neighborhood of $\Psi_n(\alpha_q(X))$ for n sufficiently large. To see this, observe that $F_n^*(y) = H(\Psi_n^{-1}(y))$, where H is the cdf of X. By ($\mathcal{A}$-4), there exists a $\delta > 0$ and open set B containing $\alpha_q(X)$ such that for all $x \in B$, $\Psi_n'(x)$ is bounded above δ for all n greater than some finite n_0. Apply the inverse function theorem to show that

$$f_n^*(\Psi_n(x)) = h(\Psi_n^{-1}(\Psi_n(x)))\Psi_n^{-1\prime}(\Psi_n(x)) = \frac{h(x)}{\Psi_n'(x)} < h(x)/\delta$$

for all $x \in B$ and $n > n_0$. Therefore, for $n > n_0$ and any $\eta > 0$, there is an $\epsilon^* > 0$ such that

$$\begin{aligned}
&\max\{F_n^*(\Psi_n(\alpha_q(X)) + \epsilon) - F_n^*(\Psi_n(\alpha_q(X))), F_n^*(\Psi_n(\alpha_q(X))) - F_n^*(\Psi_n(\alpha_q(X)) - \epsilon)\} \\
&\qquad < \epsilon\eta + \epsilon f_n^*(\Psi_n(\alpha_q(X))) < \epsilon(\eta + h(x)/\delta)
\end{aligned}$$

for all $\epsilon < \epsilon^*$. Because η and ϵ can be made arbitrarily small and $\delta > 0$, the expression

$$\max\{F_n^*(\Psi_n(\alpha_q(X)) + \epsilon) - F_n^*(\Psi_n(\alpha_q(X))), F_n^*(\Psi_n(\alpha_q(X))) - F_n^*(\Psi_n(\alpha_q(X)) - \epsilon)\}$$

can also be made arbitrarily small for all n sufficiently large.

By Proposition 1 and the dominated convergence theorem, $L_n - \mathrm{E}[L_n|X]$ converges to zero almost surely, which implies convergence in probability as well. Therefore, for any choice of $\epsilon > 0$, $\Pr(|L_n - \mathrm{E}[L_n|X]| > \epsilon)$ can be made arbitrarily small by letting $n \to \infty$.

The assumption that X is continuous implies $H(\alpha_q(X)) = q$ exactly. The regularity assumptions on the ψ_i imply that $\Psi_n(x)$ is strictly increasing around $x = \alpha_q(X)$. Therefore, $F_n^*(\mathrm{E}[L_n|\alpha_q(X)]) = q$ exactly as well. From equation (32), we conclude that $|F_n(\mathrm{E}[L_n|\alpha_q(X)]) - q|$ vanishes as $n \to \infty$. **QED**

D Proof of Proposition 6

The proof of Proposition 6 follows almost directly from two lemmas. The first provides the rate of convergence of the moments of L_n to their asymptotic values. The second shows that quantiles converge at the same rate.

Lemma 6

Let $Z_1, Z_2, \ldots,$ be a sequence of independent and identically distributed random variables, and let $Y_n \equiv (1/n)\sum_{i=1}^n Z_i$ denote the mean of the partial sums. If all moments of the Z_i exist, then $\mathrm{E}[Y_n]$=$\mathrm{E}[Z_1]$ and

$$\mathrm{E}[Y_n^j] = \mathrm{E}[Z_1]^j + O(n^{-1})$$

for $j \geq 2$.

Proof: Let $M_Z(t)$ be the moment generating function ("mgf") of the Z_i and let $M_n(t)$ be the mgf of the Y_n. Because the Z_i are independent, $M_n(t)$ is given by

$$M_n(t) = M_Z(t/n)^n = \left(\sum_{j=0}^{\infty} \frac{\mu_j}{j! n^j} t^j\right)^n, \tag{33}$$

where the μ_j are the uncentered moments of the Z_i. Using Gradshteyn and Ryzhik (1965, eq. 0.314), we can expand the exponentiated polynomial as

$$M_n(t) = \sum_{j=0}^{\infty} c_j t^j, \tag{34}$$

where the coefficients c_j are

$$\begin{aligned} c_0 &= \mu_0 = 1; \\ c_j &= \frac{1}{j}\sum_{k=1}^{j}(kn - j + k)\frac{\mu_k}{k! n^k} c_{j-k} \quad \text{for } j > 0. \end{aligned}$$

For $j = 1$, we get $c_1 = \mu_1$. For $j > 1$, we rearrange the expression for c_j to get

$$c_j = \frac{1}{j}\mu_1 c_{j-1} + \frac{1}{j}\sum_{k=1}^{j-1} \frac{\mu_{k+1} c_{j-k-1} - (j-k)\mu_k c_{j-k}}{k! n^k}. \tag{35}$$

For $j = 2$, we have $c_2 = \mu_1 c_1/2 + O(n^{-1}) = \mu_1^2/2! + O(n^{-1})$. For higher j, we use induction. If

$$c_{j-1} = \frac{\mu_1^{j-1}}{(j-1)!} + O(n^{-1}),$$

then equation (35) implies that

$$c_j = \frac{\mu_1}{j}\frac{\mu_1^{j-1}}{(j-1)!} + O(n^{-1}) = \frac{\mu_1^j}{j!} + O(n^{-1}),$$

as required. By definition of the moment generating function, the first moment of Y_n is c_1 and the j^{th} moment is $c_j j!$, which completes the proof. **QED**

Lemma 7

Let $\{Y_n\}$ be a sequence of random variables converging in distribution to Y^*, and let κ_{nj} and κ_j^* be the j^{th} cumulants of Y_n and Y^*, respectively. If the distribution of Y^* is arbitrarily differentiable, and if $(\kappa_{nj} - \kappa_j^*) = O(n^{-1})$, then

$$\alpha_q(Y_n) - \alpha_q(Y^*) = O(n^{-1}).$$

Proof is a mechanical application of a generalized Cornish-Fisher expansion due to Hill and Davis (1968) for a sequence of distributions converging to an arbitrarily differentiable limiting distribution. Their expansion can be re-arranged to show that the difference of quantiles equals

$$\sum_{j=1}^{\infty}(\kappa_{nj} - \kappa_j^*)\Xi_j(\alpha_q(Y^*))$$

plus terms depending on products of the $(\kappa_{nj} - \kappa_j^*)$, where the functions $\Xi_j(y)$ depend only on derivatives of the density of Y^*.

Proof of Proposition 6: By Lemma 6, the conditional moments of L_n are given by $\mathrm{E}[L_n|X] = \psi(X)$ and

$$\mathrm{E}[L_n^j|X] = \psi(X)^j + O(n^{-1})$$

for $j \geq 2$. By the dominated convergence theorem, the unconditional moments are $\mathrm{E}[L_n] = \mathrm{E}[\psi(X)]$ and

$$\mathrm{E}[L_n^j] = \mathrm{E}[\psi(X)^j] + O(n^{-1}) \tag{36}$$

for $j \geq 2$. Note that the $\mathrm{E}[\psi(X)^j]$ are simply the moments of the distribution of $\psi(X)$.

Let κ_{nj} and κ_j^* denote the j^{th} cumulants of L_n and $\psi(X)$, respectively. First cumulants equal first moments, so $\kappa_{n1} = \kappa_1^*$ exactly. Using standard formulae for conversion of moments to cumulants (e.g., Stuart and Ord 1994, eq. 3.40), we can easily show that equation (36) implies

$$\kappa_{nj} - \kappa_j^* = O(n^{-1})$$

for $j \geq 2$.

The distribution of $\psi(X)$ is $F^*(y) = H(\psi^{-1}(y))$. By assumption, ψ is increasing, continuous and arbitrarily differentiable, so ψ^{-1} also has these properties. As the distribution H of X is arbitrarily continuous, F^* must be arbitrarily continuous as well. The conditions of Lemma 7 are satisfied, which completes the proof. **QED**

E Proof of Lemma 1

Divide Ω into two subsets

$$\begin{aligned} B_1 &= \{\omega : 0 \leq \min(Y_1(\omega), Y_2(\omega)) \vee \max(Y_1(\omega), Y_2(\omega)) \leq 0\} \\ B_2 &= \{\omega : (Y_1(\omega) < 0 < Y_2(\omega)) \vee (Y_2(\omega) < 0 < Y_1(\omega))\}. \end{aligned}$$

Observe that $B_1 \cup B_2 = \Omega$ and $B_1 \cap B_2 = \emptyset$. If Y is an integrable random variable on $(\Omega, \mathcal{F}, P)$, we can write

$$\begin{aligned} \mathrm{E}[Y^+] &= \int_\Omega \max(Y(\omega), 0)\, P(d\omega) \\ &= \int_{B_1} \max(Y(\omega), 0)\, P(d\omega) + \int_{B_2} \max(Y(\omega), 0)\, P(d\omega). \end{aligned}$$

The set B_1 contains all ω for which Y_1 and Y_2 are either both positive or both negative. Under both these circumstances, $\max(Y_1(\omega) + Y_2(\omega), 0)$ equals $\max(Y_1(\omega), 0) + \max(Y_2(\omega), 0)$, so

$$\int_{B_1} \max(Y_1(\omega) + Y_2(\omega), 0)\, P(d\omega) = \int_{B_1} \max(Y_1(\omega), 0)\, P(d\omega) + \int_{B_1} \max(Y_2(\omega), 0)\, P(d\omega). \quad (37)$$

The set B_2 contains all ω for which Y_1 and Y_2 are of opposite sign, so

$$\int_{B_2} \max(Y_1(\omega) + Y_2(\omega), 0)\, P(d\omega) \leq \int_{B_2} \max(Y_1(\omega), 0)\, P(d\omega) + \int_{B_2} \max(Y_2(\omega), 0)\, P(d\omega). \quad (38)$$

Summing left and right hand sides of equations (37) and (38), we obtain

$$\mathrm{E}[(Y_1 + Y_2)^+] \leq \mathrm{E}[Y_1^+] + \mathrm{E}[Y_2^+]. \quad (39)$$

If $P(B_2) > 0$, then the inequality in equation (38) is strict, and therefore the inequality in equation (39) is strict as well.

F Asymptotic EEL in CreditRisk$^+$

I derive the asymptotic EEL for a homogeneous portfolio under a single systematic factor version of CreditRisk$^+$. Let $\bar{p}$ denote default probability, λ denote LGD, w denote factor loading, and σ denote the volatility of systematic factor X. The conditional expected loss rate in the CreditRisk$^+$ specification is given by equation (14). As $n \to \infty$, L_n converges to $\psi(X)$, so EEL converges to

$$\mathrm{EEL}[L_\infty, c] = \mathrm{E}[(\psi(X) - c)^+] = \int_{\psi^{-1}(c)}^{\infty} (\psi(x) - c) h(x)\, dx, \tag{40}$$

where $h(\cdot)$ is the gamma pdf with mean one, variance σ^2. Using Abramowitz and Stegun (1968, 6.5.1, 6.5.21) to solve this integral, I obtain

$$\mathrm{EEL}[L_\infty, c] = (EL - c)(1 - H(\psi^{-1}(c))) + \frac{EL \cdot w}{\Gamma(1 + 1/\sigma^2)} \left(\psi^{-1}(c)/\sigma^2\right)^{1/\sigma^2} \exp\left(-\psi^{-1}(c)/\sigma^2\right), \tag{41}$$

where $H(\cdot)$ denotes the gamma cdf, EL is expected loss ($\lambda\bar{p}$), and

$$\psi^{-1}(c) = \frac{c - (1 - w) \cdot EL}{w \cdot EL}.$$

The gamma cdf does not have neat closed form, but poses no numerical difficulties. Standard software for solving nonlinear equations quickly finds the capital ratio c which covers EEL target θ. In the special case of $\sigma = 1$, the gamma distribution reduces to the exponential distribution, and equation (41) simplifies to

$$\mathrm{EEL}[L_\infty, c] = w \cdot EL \cdot \exp\left(\psi^{-1}(c)\right).$$

To hit an EEL target of θ, we invert this equation to get

$$c = EL - w \cdot EL \cdot (1 + \ln(\theta) - \ln(w \cdot EL)).$$

References

Abramowitz, Milton and Irene A. Stegun, *Handbook of Mathematical Functions* number 55. In 'Applied Mathematics Series.', National Bureau of Standards, 1968.

Basel Committee on Bank Supervision, "Credit Risk Modelling: Current Practices and Applications," Technical Report, Bank for International Settlements 1999.

——, "The Internal Ratings-Based Approach: Supporting Document to the New Basel Capital Accord," Technical Report, Bank for International Settlements 2001.

Billingsley, Patrick, *Probability and Measure,* third ed., New York: John Wiley & Sons, 1995.

Bürgisser, Peter, Alexandre Kurth, and Armin Wagner, "Incorporating Severity Variations into Credit Risk," 2000.

Carey, Mark, "Dimensions of Credit Risk and Their Relationship to Economic Capital Requirements," in Frederic S. Mishkin, ed., *Prudential Supervision: Why Is It Important and What Are the Issues?*, NBER and UC Press, forthcoming.

Credit Suisse Financial Products, *CreditRisk+: A Credit Risk Management Framework,* London: Credit Suisse Financial Products, 1997.

Finger, Christopher C., "Conditional Approaches for CreditMetrics Portfolio Distributions," *CreditMetrics Monitor*, 1999, pp. 14–33.

Frye, Jon, "Collateral Damage: A Source of Systematic Credit Risk," *Risk*, 2000.

GARP Committee on Regulation and Supervision, "Response to Basle's "Credit Risk Modelling: Current Practices and Applications"," Technical Report, Global Association of Risk Professionals 1999.

Gordy, Michael B., "A Comparative Anatomy of Credit Risk Models," *Journal of Banking and Finance*, January 2000, *24* (1-2), 119–149.

Gradshteyn, I.S. and I.M. Ryzhik, *Table of Integrals, Series, and Products,* New York: Academic Press, 1965.

Gupton, Greg M., Christopher C. Finger, and Mickey Bhatia, *CreditMetrics–Technical Document,* New York: J.P. Morgan & Co. Incorporated, 1997.

Hill, G.W. and A.W. Davis, "Generalized Asymptotic Expansions of Cornish-Fisher Type," *Annals of Mathematical Statistics*, August 1968, *39* (4), 1264–1273.

Jones, David, "Emerging problems with the Basel Capital Accord: Regulatory capital arbitrage and related issues," *Journal of Banking and Finance*, January 2000, *24* (1-2), 35–58.

Knopp, Konrad, *Infinite Sequences and Series,* New York: Dover Publications, 1956.

McAllister, Patrick H. and John J. Mingo, "Bank capital requirements for securitized loan pools," *Journal of Banking and Finance*, September 1996, *20* (8), 1381–1405.

Petrov, Valentin V., ed., *Limit Theorems of Probability Theory* number 4. In 'Oxford Studies in Probability.', Oxford University Press, 1995.

Stuart, Alan and J. Keith Ord, *Kendall's Advanced Theory of Statistics,* sixth ed., Vol. 1, London: Oxford University Press, 1994.

Treacy, William F. and Mark S. Carey, "Credit Risk Rating at Large U.S. Banks," *Federal Reserve Bulletin*, 1998, *84* (11), 897–921.

Vasicek, Oldrich A., "The Loan Loss Distribution," Technical Report, KMV Corporation December 1997.

White, Halbert, *Asymptotic Theory for Econometricians* Economic Theory, Econometrics, and Mathematical Economics, Academic Press, 1984.