Available online at www.sciencedirect.com

SCIENCE DIRECT

Journal of Financial Intermediation 12 (2003) 199–232

Journal of Financial Intermediation

www.elsevier.com/locate/jfi

A risk-factor model foundation for ratings-based bank capital rules

Michael B. Gordy

Board of Governors of the Division of Research and Statistics, Board of Governors of the Federal Reserve System, Washington, DC 20551, USA

Received 25 October 2002

Abstract

I demonstrate that ratings-based capital rules, including both the current Basel Accord and its proposed revision, can be reconciled with the general class of credit value-at-risk models. Each exposure's contribution to VaR is *portfolio-invariant* only if (a) dependence across exposures is driven by a single systematic risk factor, and (b) no exposure accounts for more than an arbitrarily small share of total portfolio exposure. Analysis of rates of convergence to asymptotic VaR leads to a simple and accurate portfolio-level add-on charge for undiversified idiosyncratic risk. There is no similarly simple way to address violation of the single factor assumption.

JEL classification: G31; G38

Keywords: Capital allocation; Banking regulation; Value-at-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

E-mail address: mgordy@frb.gov.

1042-9573/$ – see front matter
doi:10.1016/S1042-9573(03)00040-8

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. Once allocated to the individual instruments as capital charges, economic capital provides a shadow price on the cost of holding each position. Directly or indirectly, model applications to portfolio management depend on the capacity to assign appropriate instrument-level capital charges.

Model-based assessment of capital charges 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 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 that arose in response to this incentive have undermined the effectiveness of regulatory capital rules (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 could 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 across institutions and that data constraints would prevent validation of key model parameters and assumptions. 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 offer 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. In the latest version of the Basel proposal for an Internal Ratings-Based (IRB) approach (Basel Committee on Bank Supervision, 2001), the bucketing system is required to partition instruments by internal borrower rating; by loan type (e.g., sovereign vs. corporate vs. project finance); by one or more proxies for seniority/collateral type, which determines loss severity in the event of default; and by maturity. More complex systems might further partition instruments by, for example, country and industry of borrower. Regardless of the sophistication of the bucketing scheme, capital charges are *portfolio-invariant*, i.e., the capital charge on a given instru-

[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. A very limited degree of risk-sensitivity is achieved through discounts to the standard 8% that are applied to certain special classes of lending, e.g., to OECD member governments, to other banks in OECD countries, and for residential mortgages.

ment 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 ratings-based capital rules. Throughout this paper, I will use the term "ratings-based" to refer broadly to portfolio-invariant capital allocation rules with bucketing along multiple dimensions, rather than to constrain the term to schemes in which capital depends only on a traditional univariate credit rating.

Though a ratings-based scheme may be a necessary "second-best" solution under current conditions, it is nonetheless desirable that the 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 a smoother path of evolution towards a regime based on internal models. This paper asks how a rigorous models-based calibration of ratings-based capital charges can be achieved. In particular, 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 other instruments are present in the portfolio.

If economic capital is defined within the value-at-risk (VaR) paradigm, then the problem has a simple answer. 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 qth quantile of the distribution of portfolio loss over the horizon. I show that two conditions are necessary and (with a few regularity conditions) sufficient to guarantee portfolio-invariance under VaR in risk-factor models: 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 at most a single systematic risk factor.

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. Capital charges calibrated to the asymptotic case, which assume that idiosyncratic risk is diversified away completely, must understate required capital for any given finite portfolio. To assess the magnitude of this bias, I examine the rate of convergence of VaR to its asymptotic limit. 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 extremely well, so that moderate departures from asymptotic granularity need not pose a problem in practice for ratings-based 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. A more robust risk measure is *expected shortfall* (ES), which is (loosely speaking) the expected loss conditional on being in the tail. From the perspective of an insurer of deposits (e.g., the FDIC in the US), an even more relevant risk measure 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 ES and EEL deliver portfolio-invariant capital charges for an asymptotic portfolio in a single-factor setting. Expected shortfall does, but EEL does not, and thus is unsuitable as a soundness standard for deriving risk-bucket capital charges.

The emphasis in this paper is on generality across portfolios and models. The use of asymptotics to characterize model properties is not new, but all previous analyses have been applied to homogeneous portfolios and with the objective of simplifying computation.[2] 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.

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 (with minor additional restrictions) to the case of "multi-state" models in which loss is measured on a market-value basis. A capital adjustment for undiversified idiosyncratic risk is developed in Section 4. In Section 5, I examine the asymptotic behavior of expected shortfall and expected excess loss as alternatives to VaR. 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 a book-value (or *actuarial*) 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.

[2] 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 et al. (2001) characterize the asymptotic behavior of a generalized CreditRisk$^+$ model on a sequence of portfolios with n statistically identical copies of a fixed heterogeneous portfolio.

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 PD or expected default frequency, 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.

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 unspecified. Regardless of their identity, it is assumed that all dependence across credit events is 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's Portfolio Manager. The similarity to CreditRisk$^+$ is easiest to see because that model is written in the language of conditional default probabilities (Credit Suisse Financial Products, 1977). 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 $[0, 1]$.[3]

CreditMetrics (Gupton et al., 1997), 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 $\mathbb{N}(0, \Omega)$. Associated with each obligor is a latent variable R_i which represents the return on the firm's assets. R_i is given by

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

where the ϵ_i are i.i.d. $\mathbb{N}(0, 1)$ white noise (representing obligor-specific risk). Without loss of generality, the weights w_i and ψ_i are scaled so that R_i is mean zero, variance one.[4] A borrower defaults if and only if its asset return falls below a threshold value γ_i.

[3] 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.

[4] The usual specification 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. The scaling of R_i implies that the weight on the idiosyncratic factor is given by $\psi_i = (1 - w_i' \Omega w_i)^{1/2}$.

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

$$p_i(x) = \Pr\bigl(\epsilon_i \leqslant (\gamma_i + xw_i)/\psi_i\bigr) = \Phi\bigl((\gamma_i + xw_i)/\psi_i\bigr) \tag{3}$$

where Φ is the standard normal cumulative distribution function (cdf). To calibrate the parameter γ_i, note that the unconditional probability of default is $\Phi(\gamma_i)$, so $\gamma_i = \Phi^{-1}(\bar{p}_i)$. 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 projects have failed. 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. (2001) 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.[5] 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 we permit to be negative to accommodate short positions. The usual assumption of conditional independence of defaults is extended to conditional independence of the U_i. I assume that

($\mathcal{A}$-1) the $\{U_i\}$ are bounded in the interval $[-1, 1]$ and, conditional on X, are mutually independent.

For a portfolio of n obligors, define the *portfolio loss ratio* L_n as the ratio of total losses to total portfolio exposure,[6] i.e.,

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

[5] 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.

[6] 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.

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

$$\alpha_q(Y) \equiv \inf\bigl\{y\colon\ \Pr(Y \leqslant y) \geqslant q\bigr\}. \tag{5}$$

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

Suppose that the bank selects its portfolio as the first n elements 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

($\mathcal{A}$-2) 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 $A_n / \sum_{i=1}^n A_i = O(n^{-(1/2+\zeta)})$,

where order notation $O(\cdot)$ is defined as in Billingsley (1995, A18). The restrictions in ($\mathcal{A}$-2) are sufficient to guarantee that the share of the largest single exposure in total portfolio exposure vanishes to zero as the number of exposures in the portfolio increases. As a practical matter, the restrictions are quite weak and would be 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 positive minimum size and from above by a finite maximum size.

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* ($\mathcal{A}$-1) *and* ($\mathcal{A}$-2) *hold, 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* ($\mathcal{A}$-1) *and* ($\mathcal{A}$-2) *hold, then* $\mathrm{V}[L_n] - \mathrm{V}[\mathrm{E}[L_n|X]] \to 0$.

The proof is in Appendix A.

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

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

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* ($\mathcal{A}$-1) *and* ($\mathcal{A}$-2) *hold, then for any* $\epsilon > 0$,

$$F_n\big(\alpha_q\big(\mathrm{E}[L_n|X]\big) + \epsilon\big) \to [q, 1], \tag{7}$$

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

The proof is in Appendix B. For all practical purposes, this proposition ensures that Eq. (6) will hold.[7] The literal interpretation of Proposition 3 is that if capital is strictly greater than the qth quantile of $\mathrm{E}[L_n|X]$, then it 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, if capital is strictly less than the qth quantile of $\mathrm{E}[L_n|X]$, then it is guaranteed, in the limit, to fail to cover the qth 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 may be relatively 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. Indeed, the portfolio need not be limited to traditional loans or bonds, but could include credit derivatives and guarantees, tranches of structured financial products such as collateralized debt obligations (CDOs) and asset-backed securitizations (ABSs), and so on.[8] We have bounded the support of the U_i, but have otherwise not restricted the behavior of the conditional expected loss functions (i.e., the $\mathrm{E}[U_i|x]$).[9] 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:

---

[7] 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 $n \to \infty$.

[8] For extensions of the results of this paper to determining economic capital requirements on CDO and ABS tranches, see Pykhtin and Dev (2002, 2003) and Gordy and Jones (2003).

[9] 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 this restriction to a requirement that the U_i have bounded variance. See the modified version of ($\mathcal{A}$-1) introduced in Section 3.

($\mathcal{A}$-3) the systematic risk factor X is one-dimensional; and
($\mathcal{A}$-4) there is an open interval B containing $\alpha_q(X)$ and a real number $n_0 < \infty$ such that
(i) for all i, $\mathrm{E}[U_i|x]$ is continuous in x on B,
(ii) $E[L_N|x]$ is nondecreasing in x on B for all $n > n_0$, and
(iii) $\inf_{x\in B} \mathrm{E}[L_n|x] \geqslant \sup_{x\leqslant \inf B} \mathrm{E}[L_n|x]$ and $\sup_{x\in B} \mathrm{E}[L_n|x] \leqslant \inf_{x\geqslant \sup B} \mathrm{E}[L_n|x]$ for all $n > n_0$.

Intuitively, assumption ($\mathcal{A}$-3) imposes a single global business cycle as the source of all dependence across exposures. With assumption ($\mathcal{A}$-4), it guarantees that the neighborhood of the qth quantile of $\mathrm{E}[L_n|X]$ is associated with the neighborhood of the (unique) qth quantile of X. Without ($\mathcal{A}$-4), the tail quantiles of the loss distribution would depend in complex ways on how conditional expected loss for each borrower varies with x. A more parsimonious way to avoid this problem would have been to require that the $\mathrm{E}[U_i|x]$ be nondecreasing in x for all i. However, such a requirement would exclude hedging instruments (such as credit derivatives) and obligors with counter-cyclical credit risk. Assumption ($\mathcal{A}$-4) allows for *some* U_i to be negatively associated with X, just so long as, asymptotically and in aggregate, such instruments do not alter the monotonic dependence of losses on the systematic factor when X is near the relevant "tail event." Furthermore, ($\mathcal{A}$-4) allows $\mathrm{E}[L_n|x]$ to be discontinuous or locally nonmonotonic in x when x is not in the neighborhood of $\alpha_q(X)$.

For notational convenience, define functions $\mu_i(x) \equiv \mathrm{E}[U_i|x]$ and

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

We now have

Proposition 4. *If* ($\mathcal{A}$-3) *and* ($\mathcal{A}$-4) *are satisfied, then* $\alpha_q(\mathrm{E}[L_n|X]) = \mathrm{E}[L_n \mid \alpha_q(X)] = M_n(\alpha_q(X))$ *for* $n > n_0$.

Proof. Fix $n > n_0$. If $X \leqslant \alpha_q(X)$, then $M_n(X) \leqslant M_n(\alpha_q(X))$, so

$$\Pr\bigl(M_n(X) \leqslant M_n\bigl(\alpha_q(X)\bigr)\bigr) \geqslant \Pr\bigl(X \leqslant \alpha_q(X)\bigr) \geqslant q.$$

Fix any $y < M_n(\alpha_q(X))$, and let $\hat{x} = \sup\{x\colon\ M_n(x) \leqslant y\}$. ($\mathcal{A}$-4) guarantees that $\hat{x} < \alpha_q(X)$, so

$$\Pr\bigl(M_n(X) \leqslant y\bigr) \leqslant \Pr\bigl(X < \hat{x}\bigr) < q.$$

Therefore,

$$\inf\bigl\{y\colon\ \Pr\bigl(M_n(X) \leqslant y\bigr) \geqslant q\bigr\} = M_n\bigl(\alpha_q(X)\bigr). \qquad \square$$

The importance of this result lies in the linearity of the expectations operator. Whereas $\alpha_q(\mathrm{E}[L_n|X])$ may in the general case be highly complicated, $\mathrm{E}[L_n \mid \alpha_q(X)]$ is simply the exposure-weighted average of the individual assets' conditional expected losses. Taken together with Propositions 1 and 3, Proposition 4 permits a simple and powerful rule for determining capital requirements. For asset i, allocate capital per dollar book value (inclusive of expected loss) of $c_i \equiv \mu_i(\alpha_q(X)) + \epsilon$, for some arbitrarily small ϵ.[10] Observe that this capital charge depends only on the characteristics of instrument 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,

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

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 in order to eliminate the possibility of discontinuities at the desired quantiles, the insolvency probability converges to $1-q$ exactly for $\epsilon = 0$. A simple way to achieve this would be to require that X be continuous and that the $\mu_i(x)$ functions have bounded derivatives. However, we can be rather less restrictive, as we really need only to guarantee that the asymptotic portfolio loss cdf is smooth and has bounded derivatives in the neighborhood of its qth quantile value. The following condition is sufficient to circumvent the technical caveats of Proposition 3.

($\mathcal{A}$-5) There exists an open interval B containing $\alpha_q(X)$ on which
  (i) the cdf of systematic factor X is continuous and increasing,
  (ii) for all i, $\mu_i(x)$ is differentiable, and
  (iii) there are real numbers $\underline{\delta}, \bar{\delta}$ and $n_0 < \infty$ such that $0 < \underline{\delta} \leqslant M_n'(x) \leqslant \bar{\delta} < \infty$ for all $n > n_0$.

This assumption allows for a non-trivial share of the portfolio to consist of hedged instruments or loans to counter-cyclical borrowers.

In Appendix C, I show that

Proposition 5. *If assumptions* ($\mathcal{A}$-1)–($\mathcal{A}$-5) *hold, then* $\Pr(L_n \leqslant \mathrm{E}[L_n \mid \alpha_q(X)]) \to q$ *and* $|\alpha_q(L_n) - \mathrm{E}[L_n \mid \alpha_q(X)]| \to 0$.

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

[10] 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.

The results of this section closely parallel recent developments in techniques for capital allocation in a market risk setting. Gouriéroux et al. (2000), Tasche (2000) and others show how to take partial first derivatives of VaR.[11] In terms of the notation used here, the first derivative is given by

$$\frac{\mathrm{d}\alpha_q(L_n)}{\mathrm{d}A_i} = \mathrm{E}\big[U_i \mid L_n = \alpha_q(L_n)\big]. \tag{11}$$

Under the assumptions of Proposition 5, the condition $L_n = \alpha_q(L_n)$ is asymptotically equivalent to $X = \alpha_q(X)$, which implies that marginal VaR is equal to $\mu_i(\alpha_q(X))$. Gouriéroux et al. (2000) require that the joint distribution of the losses $\{U_i\}$ be continuous, as otherwise VaR need not be differentiable. This presents a problem in application to credit risk modeling, as credit risk is largely driven by discrete events (e.g., defaults). The approach taken here in obtaining Proposition 5 allows for the distribution of U_i to be discrete or mixed.[12]

Portfolio-invariance depends strongly on the asymptotic assumption and on the assumption of a single systematic risk factor. 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. As a practical matter, residual idiosyncratic risk need not be an impediment to ratings-based capital allocation. Large internationally active banks are typically (though not invariably) near the asymptotic ideal. Furthermore, the techniques of Section 4 allow for a simple portfolio-level correction.

Assumption ($\mathcal{A}$-3) is much less innocuous from an empirical point-of-view. It can be relaxed only slightly. Say that some group of obligors shared dependence on a "local" risk factor. Conditional on the global risk factor X, the $\{U_i\}$ within the group would no longer be independent, though they would remain independent of the $\{U_j\}$ outside the group. So long as the within-group exposures *in aggregate* account for a trivial share of the total portfolio (i.e., they could be aggregated into a single exposure without violating assumption ($\mathcal{A}$-2)), the local dependence can be ignored.

Even the largest banks have geographic and industrial concentrations at some level. If these larger-scale sectors are not perfectly comonotonic, then portfolio-invariance is lost. Say we had two risk factors, and obligors could differ in their sensitivity to each factor. The realizations (x_1, x_2) associated with a given quantile of the loss distribution would then 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

[11] This problem was solved independently by several authors. See references in Tasche (2002, Section 4).

[12] Tasche (2000) provides slightly less stringent conditions for differentiability. Tasche (2001) applies Eq. (11) to a discrete model (CreditRisk$^+$), and discusses the technical issues that arise.

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 comonotonic), 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 μ_i functions need not even share a common form across instruments. Sorting 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.[13]

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 that 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.[14] 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 size" A_i as the current market value.

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 $\mu_i(x) \equiv \mathrm{E}[U_i|x]$ 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 Eq. (2)), such that obligor i defaults if $R_i \leqslant \gamma_{i,G}$, and transits to "live" grade g if $\gamma_{i,g} < R_i \leqslant \gamma_{i,g-1}$. The variables $(X, \epsilon_1, \epsilon_2, \ldots, \epsilon_n)$ are i.i.d. $\mathbb{N}(0, 1)$. Therefore, the conditional transition probabilities are given in CreditMetrics

[13] 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.

[14] 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.

by

$$p_{ig}(x) = \Phi\Big((\gamma_{i,g-1} + x w_i)\Big/\sqrt{1-w_i^2}\Big) - \Phi\Big((\gamma_{i,g} + x w_i)\Big/\sqrt{1-w_i^2}\Big), \tag{12}$$

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

Consider a zero-coupon instrument maturing at or 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

$$\mathrm{MTM}_i(x) = \sum_{g=1}^{G} v_{ig}(x) p_{ig}(x) + \bar{A}_i\big(1 - \mathrm{E}[\mathrm{LGD}_i|x]\big) p_{i,G+1}(x), \tag{13}$$

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 $\mu_i(x)$ are then given by

$$\mu_i(x) = \frac{\exp(-rT)}{A_i}\big(\mathrm{E}\big[\mathrm{MTM}_i(X)\big] - \mathrm{MTM}_i(x)\big), \tag{14}$$

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

The results of the previous section can be adapted to a mark-to-market setting without difficulty. In contrast to the actuarial case, the MTM loss rate is not necessarily bounded. To accommodate the MTM case, I modify assumption ($\mathcal{A}$-1) as follows:

($\mathcal{A}$-1) Conditional on X, the $\{U_i\}$ are independent. The conditional second moment of loss exists and is bounded; i.e., there exists a function $\Upsilon(x)$ such that $\mathrm{E}[U_i^2|x] \leqslant \Upsilon(x) < \infty$ for all instruments i and realizations x. Furthermore, $\mathrm{E}[\Upsilon(X)] < \infty$.

This version of the assumption is strictly weaker than the version of Section 1.

For a given portfolio of n assets, L_n, as defined in Eq. (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, under the relaxed version of assumption ($\mathcal{A}$-1). Indeed, the proofs given in the appendix explicitly rely only on the relaxed version. The results in no way depend on the assumptions and conventions of CreditMetrics, which are described above for illustrative purposes.[15] By the same logic as before, the appropriate asymptotic capital charge per dollar current market value for asset i is simply $\mu_i(\alpha_q(X))$.

[15] In the spirit of KMV Portfolio Manager, for example, one could replace Eq. (12) with the conditional density function for the default probability at the horizon. The summation in Eq. (13) would be replaced by an integral, and the v_{ig} would be obtained using risk-neutral valuation. Valuation in the default state in Eq. (13) also would be modified.

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 ratings-based 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.

Consider a homogeneous portfolio in which each instrument has the same conditional expected loss function $\mu(x)$ and the same exposure size. Under assumptions ($\mathcal{A}$-3) and ($\mathcal{A}$-4) and suitable regularity conditions,

$$\alpha_q(L_n) = \mu\bigl(\alpha_q(X)\bigr) + O\bigl(n^{-1}\bigr). \tag{15}$$

That is, the difference between the VaR for a given finite homogeneous portfolio and its asymptotic approximation is proportional to $1/n$.

One way to obtain this result is through a generalized Cornish–Fisher expansion due to Hill and Davis (1968) for a sequence of distributions converging to an arbitrarily differentiable limiting distribution. The jth term in the expansion of $\alpha_q(L_n)$ is proportional to the difference between the jth cumulants of the distributions for L_n and L_∞. Under very general conditions, the cumulants (for $j \geqslant 2$) converge at $O(n^{-1})$. The difficulty is in specifying precisely a set of regularity conditions under which the Cornish–Fisher expansion is guaranteed to be convergent.

Building on the results of Gouriéroux et al. (2000), Martin and Wilde (2002) derive Eq. (15) more rigorously as a Taylor series expansion of VaR around its asymptotic value. Although the necessary regularity conditions remain slightly opaque, the main additional requirement is that the conditional variance $\mathrm{V}[U|x]$ is locally continuous and differentiable in x. Furthermore, Martin and Wilde show that the $O(n^{-1})$ term is given by β/n where[16]

$$\beta = \frac{-1}{2h(x)} \frac{\mathrm{d}}{\mathrm{d}x}\left(\frac{\mathrm{V}[U|x]h(x)}{\mu'(x)}\right)\bigg|_{x=\alpha_q(X)} \tag{16}$$

and where $h(x)$ is the density of X.

Of course, Eq. (15) is in 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 "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

[16] Eq. (16) is obtained through less formal arguments in Wilde (2001).

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

$$\mu(x) = \mathrm{E}[\mathrm{LGD}] \cdot \bar{p}\bigl(1 + w(x-1)\bigr). \tag{17}$$

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 cdf of L_n can then be obtained as

$$\Pr(L_n \leqslant y) = \sum_{m=0}^{\infty} \pi_m G_m(ny). \tag{18}$$

Long before m approaches n, the π_m become negligibly small, so numerical calculation of Eq. (18) 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.

For this model, the asymptotic slope β is given by

$$\beta = \frac{1}{2\lambda}\bigl(\lambda^2 + \eta^2\bigr)\left(\frac{1}{\sigma^2}\left(1 + \frac{\sigma^2 - 1}{\alpha_q(X)}\right)\left(\alpha_q(X) + \frac{1-w}{w}\right) - 1\right). \tag{19}$$

This formula generalizes a formula derived in Wilde (2001) under the specific parameter values used in the Basel proposal.

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 asset-value correlation of 15% 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

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

Table 1
Convergence of value-at-risk

			$\mathrm{VaR}_q[L_n]$ for values of n					
Rating	$\bar{p}$	w	200	500	1000	2000	5000	∞
A	0.06	1.011	0.723	0.521	0.445	0.406	0.381	0.364
BBB	0.20	0.836	1.425	1.190	1.106	1.064	1.038	1.020
BB	1.25	0.602	5.217	4.947	4.856	4.810	4.783	4.764
B	6.25	0.415	17.881	17.584	17.485	17.435	17.405	17.385
CCC	17.50	0.295	37.663	37.335	37.226	37.172	37.139	37.117

Notes. Default probabilities and VaR expressed in percentage points. Simulations assume $q = 0.995$, $\sigma = 2$, $\lambda = 0.5$, and $\eta = 0.25$.

the value from CreditMetrics.[18] The remainder of the calibration exercise is straightforward. I choose stylized values for the default probabilities, and assume that LGD has mean 0.5 and standard deviation 0.25. The chosen coverage target is $q = 0.995$ of the loss distribution.

Results are shown in Table 1 for five rating grades. The final column ($n = \infty$) provides the asymptotic capital charge, so the difference between each column and the final column represents the "true" 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). However, the add-ons can be large relative to the asymptotic capital charge for investment grade obligors. For a homogeneous portfolio of 200 A-rated loans, the granularity add-on is roughly equal to the asymptotic charge.

Figure 1 demonstrates the relationship between the theoretical granularity add-on and $1/n$ for three homogeneous portfolios. For an extremely low-quality portfolio (CCC rating), the predicted linear relationship holds down to $n = 200$.[19] For the medium-quality (BB rated) portfolio, there are visible but negligible departures from the predicted linear relationship when $n < 500$. For a high-quality portfolio (A-rated), departures from linearity are visible at $n = 1000$ and become significant at lower values of n. Because departures from linearity are in the concave direction, a granularity adjustment calibrated to the asymptotic slope β would slightly overshoot the theoretically optimal add-on for smaller high-quality portfolios.

In the case of a non-homogeneous portfolio, determining an appropriate granularity add-on is only slightly more complex. The method of Wilde (2001) accommodates heterogeneity (the $\mathrm{V}[U|x]$ and $\mu(x)$ terms in Eq. (16) become $\mathrm{V}[L_n|x]$ and $M_n(x)$, respectively). An alternative two-step method also appears to work quite well and may be better suited to a regulatory setting. 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.

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.

19 The slope between each plotted point is constant to six significant digits for both B (not shown) and CCC portfolios.

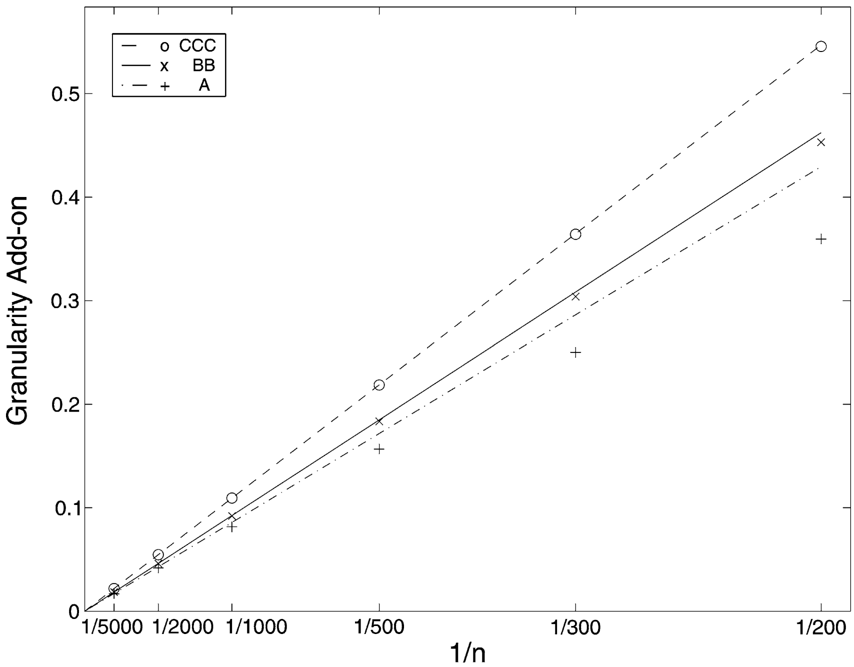

Fig. 1. Granularity add-on as linear function of $1/n$. *Note*. The "true" granularity add-on in the extended CreditRisk^{+} model is plotted with symbols for three homogeneous portfolios of various sizes. The lines show the corresponding theoretical add-on for this model.

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]

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

The higher is H_b, the more concentrated is the exposure within the bucket, so the more slowly idiosyncratic risk is diversified away. The matching methodology takes bucket-level inputs $\{\bar{p}_b, w_b, \lambda_b, \eta_b, H_b\}$ and total bucket exposure. This data structure may be especially convenient in a regulatory setting with bucket-level reporting requirements.

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

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

choose moments with intuitive interpretation. Appendix D develops a matching procedure based on five moments:[21]

- exposure-weighted expected default rate,
- expected portfolio loss rate,
- contribution of systematic risk to loss variance,
- contribution of idiosyncratic default risk to loss variance, and
- contribution of idiosyncratic recovery risk to loss variance.

Under this methodology, each of the parameters of the comparable portfolio is given by an explicit linear equation that can be interpreted as a weighted average of the characteristics of the heterogeneous portfolio. Most interestingly, the number of loans n^* can be interpreted as an inverse measure of weighted exposure concentration. Finally, the asymptotic slope β^* for the comparable homogeneous portfolio is given by Eq. (19).

The portfolio data needed for the mapping 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. Factor loadings and LGD volatilities would likely be assigned as functions of the $\bar{p}_b$ and λ_b; this is indeed the case in the proposed IRB approach (Basel Committee on Bank Supervision, 2001). The only new required inputs, the within-bucket Herfindahl indices, are easily calculated from the individual exposure sizes.

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 600 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 0.3, 0.2, 0.6, and 0.5, respectively, and the LGD volatility is set to $\eta_b = 0.5\sqrt{\lambda_b(1-\lambda_b)}$. Table 2 displays the bucket-level parameters.

Table 2
Bucket-level parameters of stylized portfolio

	$\bar{p}$	w	λ	η
1	0.05	1.040	0.3	0.229
2	0.50	0.715	0.2	0.200
3	1.00	0.629	0.6	0.245
4	5.00	0.440	0.5	0.250

Note. Default probabilities in percentage points.

[21] The matching procedure specified in the proposed granularity adjustment of Basel Committee on Bank Supervision (2001, Chapter 8) is based on a different set of moments but follows similar intuition.

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 (2001) 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 (see Treacy and Carey, 1998, Chart 3).

I first obtain parameters for the comparable homogeneous portfolio. The comparable portfolio has $n^* = 218.7$ obligors, which is under 40% of the obligor count of the original portfolio.[22] Each obligor has PD of 1.64% and factor loading $w^* = 0.487$. Loss given default has expected value 0.491 and volatility 0.247.[23] By construction, the comparable portfolio matches the original portfolio in its expected loss rate of 0.804%. For each portfolio, the standard deviation of the loss rate is 0.918%. Once the comparable portfolio is determined, the asymptotic slope β is calculated for the target quantile q. 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 the cdf in Eq. (18). 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. Despite the relatively small obligor count in the comparable portfolio, the error due to linear approximation of the "$1/n$ rule" is minimal. At $q = 99.5\%$, our approximated VaR overshoots its target by 2.2 basis points.

Table 3
Direct and approximated estimates of VaR

	Estimates	Target quantile, q		
		99.0	99.5	99.9
(i)	"True" VaR	4.577	5.522	7.872
(ii)	Asymptotic VaR	4.220	5.109	7.260
(iii)	VaR for comparable portfolio	4.570	5.535	7.872
(iv)	Granularity add-on	0.357	0.435	0.627
(v)	Approximated VaR	4.578	5.544	7.886
(vi)	Tracking error	0.001	0.022	0.014

Notes. All quantities expressed in percentage points. "True" VaR estimated by simulation with 300,000 Monte Carlo trials.

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

[23] In practice, as in this exercise, LGD volatility is often assumed to be a simple function of expected LGD. Reporting requirements and computations could be simplified by setting $\eta^* = \mathrm{VLGD}(\lambda^*)$. While this ignores the effect of nonlinearity in VLGD, the difference is typically small. In our example, $\mathrm{VLGD}(\lambda^*)$ would have been 0.250.

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 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 alternative risk measures

Industry application of credit risk modeling to capital allocation appears almost invariably to equate soundness with a coverage target for value-at-risk. However, because it ignores the distribution of losses beyond the target quantile, VaR has significant theoretical and practical shortcomings. As has been emphasized in recent literature on risk measures, VaR is not sub-additive. That is, if L_A and L_B are losses on bank portfolios A and B, then we need not have $\mathrm{VaR}_q[L_A + L_B] \leqslant \mathrm{VaR}_q[L_A] + \mathrm{VaR}_q[L_B]$, which would imply that a merger of bank A and bank B could *increase* VaR; see Frey and McNeil (2002, Section 2.3) for an example based on credit risk measurement. Sub-additivity is one of the four requirements for a "coherent" risk measure as set forth by Artzner et al. (1999).

Under the assumptions needed to achieve portfolio-invariance, $\mathrm{VaR}_q[L_A + L_B] = \mathrm{VaR}_q[L_A] + \mathrm{VaR}_q[L_B]$, so sub-additivity is preserved in mergers of asymptotically fine-grained portfolios. Even in this case, however, VaR can be manipulated by splitting the portfolio. Under an actuarial definition of loss, for example, a portfolio consisting of a single loan with default probability under $1 - q$ has $\mathrm{VaR}_q = 0$, but the same loan has positive contribution (assuming positive dependence of U on X) to VaR in an asymptotically fine-grained portfolio. Segregating this loan from the larger portfolio does not change the VaR contributions of the remaining loans, so VaR is unambiguously reduced.

Another problem is that a mean-preserving spread of the loss distribution can decrease VaR. This results in a counterintuitive non-monotonic relationship between within-portfolio correlation and VaR. Correlations increase with factor loadings, and, when factor loadings are low to moderate, VaR does as well. However, as factor loadings are pushed higher and higher, the loss distribution becomes increasingly long-tailed. VaR then shrinks towards the median, while the probability of a cataclysmic loss increases. In the limiting case of perfectly comonotonic losses, the default rate is either zero (with probability $1 - \bar{p}$) or one (with probability $\bar{p}$). If $\bar{p} < 1 - q$, then $\mathrm{VaR}_q = 0$. For a survey discussion of the potential pitfalls of VaR, see Szegö (2002).

As an alternative to VaR, Acerbi and Tasche (2002) propose using generalized *expected shortfall* (ES), defined by

$$\mathrm{ES}_q[Y] = (1 - q)^{-1}\big(\mathrm{E}[Y \cdot 1_{\{Y \geqslant \alpha_q(Y)\}}] + \alpha_q(Y)\big(q - \Pr\big(Y < \alpha_q(Y)\big)\big)\big). \tag{20}$$

The first term is often used as the definition of expected shortfall for continuous variables. It is also known as "tail conditional expectations." The second term is a correction for mass at the qth quantile of Y. Under this definition, Acerbi and Tasche (2002) show that ES is coherent and equivalent to Rockafellar and Uryasev's (2002) CVaR.

Expected shortfall offers some important advantages as a soundness standard. By Acerbi and Tasche (2002, Corollary 3.3), ES_q is continuous and monotonic in q, so small increases in the "stringency" of the capital rule (as controlled by q) lead to small increases in required capital. ES is nondecreasing with any mean-preserving spread in the loss distribution, so ES cannot decrease as factor loadings rise.

In Appendix E, I show that

Proposition 6. *If assumptions* $(\mathcal{A}\text{-}1)$–$(\mathcal{A}\text{-}5)$ *hold, then* $|\mathrm{ES}_q[L_n] - \mathrm{ES}_q[M_n(X)]| \to 0$.

An immediate implication is that ES-based capital charges are portfolio-invariant under the same assumptions as VaR-based capital charges. The asymptotic expected shortfall, $\mathrm{ES}_q[M_n(X)]$, can be decomposed as $(\sum_i c_i A_i)/(\sum_i A_i)$, where the capital charge per dollar of exposure to i is $c_i = \mathrm{E}[U_i | X \geqslant \alpha_q(X)]$. Observe that c_i depends only on how U_i depends on X, and so is portfolio-invariant. Under a variety of model specifications, c_i has an analytical solution, so there is no operational difficulty in calibrating ratings-based capital charges to an ES soundness standard.

Another alternative to VaR is *expected excess loss* (EEL). For a random variable Y and target loss rate $\theta > 0$, EEL is defined by

$$\mathrm{EEL}_\theta[Y] \equiv \inf\bigl\{y\colon \mathrm{E}\bigl[(Y-y)^+\bigr] \leqslant \theta\bigr\}, \tag{21}$$

where Y^+ denotes $\max(Y, 0)$. 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 the target loss rate. That is, the required total capital is given by $\mathrm{EEL}_\theta[L_n]$ per dollar of total exposure.

EEL is sensitive to the tail of the loss distribution, so shares many of the advantages of ES. More importantly, the target rate θ represents the expected loss borne by the depository insurance agency (such as the FDIC in the US), so EEL-based capital has a natural policy interpretation. Unlike ES, however, EEL cannot be reconciled with portfolio-invariant capital charges. 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 $\mu_a(x)$ denote the expected loss for bucket a loans conditional on $X = x$. By reasoning very similar to that of Proposition 6, we can show that

$$\mathrm{E}\bigl[(L_a - y)^+\bigr] \to \mathrm{E}\bigl[\bigl(\mu_a(X) - y\bigr)^+\bigr]$$

for any y. Therefore, the asymptotic EEL capital charge c_a is set so that $\mathrm{E}[(\mu_a(X) - c_a)^+]$ equals 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 capital for the mixed portfolio is given by $\mathrm{EEL}_\theta[\mu_m(X)]$. By construction of the mixed portfolio, we have $\mu_m(X) = (\mu_a(X) + \mu_b(X))/2$. If asymptotic EEL were portfolio-invariant, then $c_m \equiv (c_a + c_b)/2$

would satisfy

$$\theta = \mathrm{E}\big[\big(\mu_m(X) - c_m\big)^+\big]. \tag{22}$$

We now require the following triangle inequality:

Lemma 1. *If Y_1 and Y_2 are integrable random variables on a common probability space, then*

$$\mathrm{E}\big[(Y_1 + Y_2)^+\big] \leqslant \mathrm{E}\big[Y_1^+\big] + \mathrm{E}\big[Y_2^+\big].$$

If $\Pr((Y_1 < 0 < Y_2) \vee (Y_2 < 0 < Y_1)) > 0$, *then the inequality is strict.*

Proof is given in Appendix F.

The conditions of Lemma 1 apply to $Y_j \equiv (\mu_j(X) - c_j)/2$, which gives us

$$\mathrm{E}\big[\big(\mu_m(X) - c_m\big)^+\big] \leqslant \mathrm{E}\big[\big(\big(\mu_a(X) - c_a\big)/2\big)^+\big] + \mathrm{E}\big[\big(\big(\mu_b(X) - c_b\big)/2\big)^+\big] = \theta. \tag{23}$$

In general, the threshold realization of X at which $\mu_a(x) = c_a$ does not equal the corresponding threshold for portfolio b, so for some interval of x values we will have either $\mu_a(x) - c_a < 0 < \mu_b(x) - c_b$ or $\mu_a(x) - c_a > 0 > \mu_b(x) - c_b$. Therefore, the inequality in Eq. (23) 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 G, 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 (i.e., between BBB and BB). In this example, the EEL standard produces lower (higher) capital requirements than VaR for the higher (lower) grades.

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

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

Note. Capital in percentage points.

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

Note. Capital charges expressed in percentage points.

to show that the conditional expected loss rate for a mixed portfolio is

$$\mu_m(x) = \frac{1}{2}\mu_a(x) + \frac{1}{2}\mu_b(x) = \lambda \bar{p}_m \bigl(1 + w_m(x-1)\bigr) \tag{24}$$

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 $\mu_m(x)$ take on the same form as for homogeneous portfolios, so the tools of Appendix G apply without modification. Results for four different mixed portfolios are presented in Table 5. The third column shows the EEL for the mixed portfolio, while the fourth column shows the average of the EELs 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).

Discussion

Ratings-based assignment of capital charges offers significant advantages in regulatory application. The current Accord is itself a simple ratings-based framework. The proposed new Accord will 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, ratings-based systems are relatively simple to administer and need not impose burdensome reporting requirements. Validation problems are also limited in scope. As the new Accord is currently envisioned, 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 ratings-based rules 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 that could be 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.

This paper shows how risk-factor models of credit value-at-risk can be used to justify and calibrate a ratings-based system for assigning capital charges for credit risk at the instrument level. Ratings-based 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 robust 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 business cycle on all obligors. A revised Basel Accord must apply to the largest international banks, so the 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 and to prices of production inputs. 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, local events in, for example, France are permitted to contribute nothing to the default rate of French 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 judgment 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 ratings-based 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, such capacity 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.[24] In the near- to medium-term, therefore, the implausibility of the single-factor assumption need not present an obstacle

[24] 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.

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.

Acknowledgments

The views expressed herein are my own and do not necessarily reflect those of the Board of Governors or its staff. I thank Paul Calem, Darrell Duffie, Paul Embrechts, Jon Faust, Erik Heitfield, David Jones, David Lando, Gennady Samorodnistky, Dirk Tasche, Greg Udell, and Tom Wilde for their helpful comments, and Susan Yeh for editorial suggestions.

Appendix A. Proof of Propositions 1 and 2

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

Lemma 2 (Petrov, 1995, Theorem 6.7). *If $a_n \uparrow \infty$ and $\sum_{n=1}^{\infty}(\mathrm{V}[Y_n]/a_n^2) < \infty$, then*

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

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).

We now prove Proposition 1. Let $Y_n \equiv U_n A_n$ and $a_n \equiv \sum_{i=1}^{n} A_i$. For any realization x, conditional independence implies

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

Under the actuarial definition of loss, $|U_n|$ is bounded in $[0,1]$, so we must have that $\mathrm{V}[U_n|x] < 1$ for any $X = x$. For this proposition to hold under the mark-to-market paradigm as well, assumption ($\mathcal{A}$-1) provides a bound on $\mathrm{V}[U_n|x]$. Therefore, under either definition of loss, there exists a finite constant V^* such that

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

By part (b) of assumption ($\mathcal{A}$-2), the sequence $\{A_n / \sum_{i=1}^{n} A_i\}$ is $O(n^{-(1/2+\zeta)})$ for some $\zeta > 0$, so the sequence $\{(A_n / \sum_{i=1}^{n} A_i)^2\}$ is $O(n^{-(1+2\zeta)})$. By Lemma 3, the series sum must be finite. By part (a) of assumption ($\mathcal{A}$-2), 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.

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 $d_n \to 0$. Then $(1/a_n)\sum_{i=1}^n b_i d_i \to 0$.*

This result is a special case of Petrov (1995, Lemma 6.10). If we let $b_n = A_n$ and $d_n = A_n/\sum_{i=1}^n A_i$, then assumption ($\mathcal{A}$-2) guarantees that $a_n \uparrow \infty$ and $d_n \to 0$, so apply Lemma 4 to get

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

The standard rule for conditional variance gives

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

$\mathrm{E}[\mathrm{V}[U_i|X]]$ must be less than one under the actuarial paradigm. Under the mark-to-market paradigm, we have $\mathrm{E}[\mathrm{V}[U_i|X]] < \mathrm{E}[\Upsilon(X)] < \infty$ by assumption ($\mathcal{A}$-1). Therefore, there exists a finite constant V^* such that

$$\begin{aligned}\mathrm{E}\big[\mathrm{V}[L_n|X]\big] &\leqslant V^*\frac{\sum_{i=1}^n A_i^2}{(\sum_{i=1}^n A_i)^2} = V^*\frac{1}{\sum_{i=1}^n A_i}\sum_{i=1}^n \frac{A_i^2}{\sum_{j=1}^n A_j}\\ &\leqslant V^*\frac{1}{\sum_{i=1}^n A_i}\sum_{i=1}^n \frac{A_i^2}{\sum_{j=1}^i A_j} \to 0.\end{aligned}$$

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

Appendix B. Proof of Proposition 3

Almost sure convergence implies convergence in probability (see Billingsley, 1995, Theorem 25.2), so for all x and $\epsilon > 0$,

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

If F_n is the cdf of L_n, then Eq. (B.1) implies

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

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

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 qth quantile value, i.e.,

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

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

By the usual rules for conditional probability, we have

$$\begin{aligned} F_n\big(\alpha_q(\mathrm{E}[L_n|X]) + \epsilon\big) &= F_n\big(\alpha_q(\mathrm{E}[L_n|X]) + \epsilon \mid X \in S_n^+\big)\Pr\big(X \in S_n^+\big) \\ &\quad + F_n\big(\alpha_q(\mathrm{E}[L_n|X]) + \epsilon \mid X \notin S_n^+\big)\Pr\big(X \notin S_n^+\big) \\ &\geqslant F_n\big(\alpha_q(\mathrm{E}[L_n|X]) + \epsilon \mid X \in S_n^+\big)\Pr\big(X \in S_n^+\big) \\ &\geqslant F_n\big(\alpha_q(\mathrm{E}[L_n|X]) + \epsilon \mid X \in S_n^+\big)q. \end{aligned} \tag{B.2}$$

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

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

so the dominated convergence theorem (Billingsley, 1995, Theorem 16.4) implies that

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

so from Eq. (B.2) we have

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

as required.

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

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

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

$$\begin{aligned} F_n\big(\alpha_q(\mathrm{E}[L_n|X]) - \epsilon\big) &= F_n\big(\alpha_q(\mathrm{E}[L_n|X]) - \epsilon \mid X \notin S_n^-\big)\Pr\big(X \notin S_n^-\big) \\ &\quad + F_n\big(\alpha_q(\mathrm{E}[L_n|X]) - \epsilon \mid X \in S_n^-\big)\Pr\big(X \in S_n^-\big) \\ &\leqslant q + F_n\big(\alpha_q(\mathrm{E}[L_n|X]) - \epsilon \mid X \in S_n^-\big)\Pr\big(X \in S_n^-\big). \end{aligned} \tag{B.3}$$

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

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

so the dominated convergence theorem implies that

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

so from Eq. (B.3) we have

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

as required.

Appendix 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 $\epsilon > 0$,*

$$\left|F_1(y) - F_2(y)\right| \leqslant \Pr\bigl(|Y_1 - Y_2| > \epsilon\bigr) + \max\bigl\{F_2(y+\epsilon) - F_2(y), F_2(y) - F_2(y-\epsilon)\bigr\}.$$

Proof. Corollary of Petrov (1995, Lemma 1.8). □

To apply Lemma 5, let $Y_1 = L_n$ with cdf F_n and $Y_2 = \mathrm{E}[L_n|X]$ with cdf F_n^*. Fix an open set B and real numbers n_0 and $\underline{\delta}, \overline{\delta}$ for which assumptions ($\mathcal{A}$-4) and ($\mathcal{A}$-5) are satisfied. At every point $\hat{x} \in B$, we have

$$\begin{aligned}&\left|F_n\bigl(\mathrm{E}[L_n|\hat{x}]\bigr) - F_n^*\bigl(\mathrm{E}[L_n|\hat{x}]\bigr)\right| \\ &\quad\leqslant \Pr\bigl(\left|L_n - \mathrm{E}[L_n|X]\right| > \epsilon\bigr) \\ &\qquad+ \max\bigl\{F_n^*\bigl(\mathrm{E}[L_n|\hat{x}] + \epsilon\bigr) - F_n^*\bigl(\mathrm{E}[L_n|\hat{x}]\bigr), F_n^*\bigl(\mathrm{E}[L_n|\hat{x}]\bigr) - F_n^*\bigl(\mathrm{E}[L_n|\hat{x}] - \epsilon\bigr)\bigr\}\end{aligned} \tag{C.1}$$

for any $\epsilon > 0$.

For $n > n_0$, ($\mathcal{A}$-5) guarantees that $M_n(x)$ is strictly increasing on B, so for all $x \in B$, $M_n(X) \leqslant M_n(x)$ if and only if $X \leqslant x$, so $F_n^*(M_n(x)) = H(x)$ where H is the cdf of X.

Fix $\epsilon^* > 0$ such that $(\hat{x} - \epsilon^*, \hat{x} + \epsilon^*) \subset B$. For any positive $\epsilon < \underline{\delta}\epsilon^*$, we then have $\hat{x} + \epsilon/\overline{\delta} \in B$, so $M_n(\hat{x} + \epsilon/\underline{\delta}) - M_n(\hat{x}) > \epsilon$ for all $n > n_0$. As F_n^* is nondecreasing, this implies that

$$F_n^*\bigl(\mathrm{E}[L_n|\hat{x}] + \epsilon\bigr) \leqslant F_n^*\bigl(M_n(\hat{x} + \epsilon/\underline{\delta})\bigr) = H(\hat{x} + \epsilon/\underline{\delta}).$$

Similarly, we have

$$F_n^*\bigl(\mathrm{E}[L_n|\hat{x}] - \epsilon\bigr) \geqslant F_n^*\bigl(M_n(\hat{x} - \epsilon/\underline{\delta})\bigr) = H(\hat{x} - \epsilon/\underline{\delta}).$$

Thus, for all $n > n_0$,

$$\begin{aligned}&\max\bigl\{F_n^*\bigl(\mathrm{E}[L_n|\hat{x}] + \epsilon\bigr) - F_n^*\bigl(\mathrm{E}[L_n|\hat{x}]\bigr), F_n^*\bigl(\mathrm{E}[L_n|\hat{x}]\bigr) - F_n^*\bigl(\mathrm{E}[L_n|\hat{x}] - \epsilon\bigr)\bigr\} \\ &\quad\leqslant \max\bigl\{H(\hat{x} + \epsilon/\underline{\delta}) - H(\hat{x}), H(\hat{x}) - H(\hat{x} - \epsilon/\underline{\delta})\bigr\}.\end{aligned}$$

Assumption ($\mathcal{A}$-5) also provides that H is continuous and increasing on B, so for any $\eta > 0$ there exists $\epsilon > 0$ such that

$$\max\bigl\{H(\hat{x} + \epsilon/\underline{\delta}) - H(\hat{x}), H(\hat{x}) - H(\hat{x} - \epsilon/\underline{\delta})\bigr\} < \eta. \tag{C.2}$$

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 $\epsilon > 0$ and $\eta > 0$, there exists $n_\epsilon < \infty$ such that

$$\Pr\bigl(\left|L_n - \mathrm{E}[L_n|X]\right| > \epsilon\bigr) < \eta \quad \forall n > n_\epsilon. \tag{C.3}$$

Combining these results, we have that for any $\eta > 0$, there exists an $\epsilon > 0$ such that Eqs. (C.2) and (C.3) are simultaneously satisfied for $n > \max\{n_0, n_\epsilon\}$. Thus,

$$\lim_{n\to\infty} \left| F_n\big(M_n(\hat{x})\big) - F_n^*\big(M_n(\hat{x})\big)\right| \to 0. \tag{C.4}$$

Setting $\hat{x} = \alpha_q(X)$ and observing that $F_n^*(M_n(\alpha_q(X))) = H(\alpha_q(X)) = q$ establishes the first result of Proposition 5.

For any positive $\eta < \bar{\delta}\epsilon^*$ and $n > n_0$, ($\mathcal{A}$-5) implies that $M_n(\alpha_q(X)) - M_n(\alpha_q(X) - \eta/\bar{\delta}) \leqslant \eta$ so $F_n(M_n(\alpha_q(X)) - \eta) \leqslant F_n(M_n(\alpha_q(X) - \eta/\bar{\delta}))$. Because $\alpha_q(X) - \eta/\bar{\delta} \in B$, we have by Eq. (C.4) that

$$\left| F_n\big(M_n\big(\alpha_q(X) - \eta/\bar{\delta}\big)\big) - F_n^*\big(M_n\big(\alpha_q(X) - \eta/\bar{\delta}\big)\big)\right| \to 0.$$

For $n > n_0$, $F_n^*(M_n(\alpha_q(X) - \eta/\bar{\delta})) = H(\alpha_q(X) - \eta/\bar{\delta}) < q$, so there exists $\tilde{n}_- < \infty$ such that $F_n(M_n(\alpha_q(X) - \eta/\bar{\delta})) < q$ for all $n > \tilde{n}_-$. Thus, for all $n > \tilde{n}_-$, we have that $F_n(M_n(\alpha_q(X)) - \eta) < q$, which implies that $M_n(\alpha_q(X)) - \eta < \alpha_q(L_n)$.

By a parallel argument, for all positive $\eta < \bar{\delta}\epsilon^*$ there exists $\tilde{n}_+ < \infty$ such that $M_n(\alpha_q(X)) + \eta > \alpha_q(L_n)$ for all $n > \tilde{n}_+$. Thus, for all $n > \max\{\tilde{n}_-, \tilde{n}_+\}$, we have $|\alpha_q(L_n) - M_n(\alpha_q(X))| < \eta$. As η can be made arbitrarily close to zero, the second result of Proposition 5 is established.

Appendix D. Construction of the comparable homogeneous portfolio

Moment restrictions provide a convenient and intuitive way to map the heterogeneous portfolio into a homogeneous portfolio of n^* equal-sized facilities with common PD $\bar{p}^*$, factor loading w^*, and LGD parameters λ^* and η^*.

Let s_b denote the share of total portfolio exposure held in bucket b, i.e.,

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

The first two restrictions equate exposure-weighted expected default rate and expected portfolio loss rate:

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

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

$$\lambda^* = \frac{\sum_{b=1}^{B} \lambda_b \bar{p}_b s_b}{\sum_{b=1}^{B} \bar{p}_b s_b}. \tag{D.2}$$

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

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

which implies

$$w^* = \frac{\sum_{b=1}^{B} \lambda_b \bar{p}_b w_b s_b}{\sum_{b=1}^{B} \lambda_b \bar{p}_b s_b}. \tag{D.3}$$

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

$$\mathrm{E}\big[\mathrm{V}[L_n|X]\big] = \sum_{b=1}^{B} \big(\lambda_b^2\big(\bar{p}_b(1-\bar{p}_b) - (\bar{p}_b w_b \sigma)^2\big) + \bar{p}_b \eta_b^2\big) H_b s_b^2,$$

$$\mathrm{E}\big[\mathrm{V}[L^*|X]\big] = \frac{1}{n^*}\big(\lambda^{*2}\big(\bar{p}^*(1-\bar{p}^*) - (\bar{p}^* w^* \sigma)^2\big) + \bar{p}^* \eta^{*2}\big).$$

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

$$n^* = \left(\sum_{b=1}^{B} \Lambda_b H_b s_b^2\right)^{-1} \tag{D.4}$$

where

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

Finally, the variance of LGD for the comparable portfolio is given by

$$\eta^{*2} = \frac{n^*}{\bar{p}^*} \sum_{b=1}^{B} \eta_b^2 \bar{p}_b H_b s_b^2. \tag{D.5}$$

Appendix E. Proof of Proposition 6

Expected shortfall is the sum of expected loss in the tail and a correction term for mass at the VaR boundary. Under the given assumptions, the latter term disappears asymptotically both in $\mathrm{ES}_q[L_n]$ and $\mathrm{ES}_q[M_n(X)]$. For $n > n_0$, the variable $M_n(X)$ is continuous in the neighborhood of $M_n(\alpha_q(X))$ and, by Proposition 4, $\alpha_q(M_n(X)) = M_n(\alpha_q(X))$; this implies $\Pr(M_n(X) < M_n(\alpha_q(X))) = q$. By Chebyshev's inequality and assumption ($\mathcal{A}$-1), we have

$$M_n\big(\alpha_q(X)\big)^2 \leqslant \frac{\mathrm{E}[M_n(X)^2]}{\Pr(M_n(X) > M_n(\alpha_q(X)))} \leqslant \frac{\mathrm{E}[\Upsilon(X)]}{1-q} < \infty,$$

so the sequence $\{M_n(\alpha_q(X))\}$ is bounded from above. Therefore,

$$\alpha_q\big(M_n(X)\big)\big(q - \Pr\big(M_n(X) < \alpha_q\big(M_n(X)\big)\big)\big) = 0 \quad \forall n > n_0.$$

Although L_n need not be continuous, arguments parallel to those used in the proof of Proposition 5 show that $\Pr(L_n \geqslant \alpha_q(L_n)) \to 1-q$. That proposition also provides that $|\alpha_q(L_n) - M_n(\alpha_q(X))| \to 0$, so $\alpha_q(L_n)$ is asymptotically bounded from above. This implies

$$\lim_{n\to\infty} \left|\alpha_q(L_n)\bigl(q - \Pr\bigl(L_n < \alpha_q(L_n)\bigr)\bigr)\right| = 0.$$

To complete the proof, we now only need show that

$$\left|\mathrm{E}[L_n \cdot 1_{\{L_n \geqslant \alpha_q(L_n)\}}] - \mathrm{E}\bigl[M_n(X)\cdot 1_{\{M_n(X)\geqslant \alpha_q(M_n(X))\}}\bigr]\right| \to 0. \tag{E.1}$$

Let $Y_n \equiv (L_n - \alpha_q(L_n))$ and let $\widehat{Y}_n \equiv (M_n(X) - M_n(\alpha_q(X)))$. The terms of Eq. (E.1) can be re-written as

$$\begin{aligned}
&\mathrm{E}\bigl[M_n(X)\cdot 1_{\{M_n(X)\geqslant \alpha_q(M_n(X))\}}\bigr] \\
&\quad = \mathrm{E}\bigl[\bigl(M_n(X) - M_n\bigl(\alpha_q(X)\bigr)\bigr)\cdot 1_{\{M_n(X) - M_n(\alpha_q(X))\geqslant 0\}}\bigr] \\
&\qquad + M_n\bigl(\alpha_q(X)\bigr)\mathrm{E}\bigl[1_{\{M_n(X)\geqslant M_n(\alpha_q(X))\}}\bigr] \\
&\quad = \mathrm{E}\bigl[\max\bigl\{\widehat{Y}_n, 0\bigr\}\bigr] + M_n\bigl(\alpha_q(X)\bigr)\Pr\bigl(M_n(X)\geqslant M_n\bigl(\alpha_q(X)\bigr)\bigr)
\end{aligned}$$

and similarly

$$\mathrm{E}\bigl[L_n\cdot 1_{\{L_n\geqslant \alpha_q(L_n)\}}\bigr] = \mathrm{E}\bigl[\max\{Y_n, 0\}\bigr] + \alpha_q(L_n)\Pr\bigl(L_n \geqslant \alpha_q(L_n)\bigr).$$

As $|\alpha_q(L_n) - M_n(\alpha_q(X))| \to 0$ and

$$\lim_{n\to\infty}\Pr\bigl(L_n\geqslant \alpha_q(L_n)\bigr) = \lim_{n\to\infty}\Pr\bigl(M_n(X)\geqslant M_n\bigl(\alpha_q(X)\bigr)\bigr) = 1-q,$$

we have

$$\left|\alpha_q(L_n)\Pr\bigl(L_n\geqslant\alpha_q(L_n)\bigr) - M_n\bigl(\alpha_q(X)\bigr)\Pr\bigl(M_n(X)\geqslant M_n\bigl(\alpha_q(X)\bigr)\bigr)\right| \to 0.$$

For all $y, \hat{y} \in \Re$, $|\max(y,0) - \max(\hat{y},0)| \leqslant |y-\hat{y}|$. Therefore,

$$\begin{aligned}
&\left|\mathrm{E}\bigl[\max\{Y_n,0\}\bigr] - \mathrm{E}\bigl[\max\bigl\{\widehat{Y}_n,0\bigr\}\bigr]\right| \\
&\quad \leqslant \mathrm{E}\bigl[\bigl|Y_n - \widehat{Y}_n\bigr|\bigr] \leqslant \mathrm{E}\bigl[\bigl|L_n - M_n(X)\bigr|\bigr] + \bigl|\alpha_q(L_n) - M_n\bigl(\alpha_q(X)\bigr)\bigr|.
\end{aligned}$$

As each of these terms converges to zero, Eq. (E.1) is established.

Appendix F. Proof of Lemma 1

Let $(\Omega, \mathcal{F}, P)$ be the probability space for Y_1 and Y_2. Divide Ω into two subsets

$$\begin{aligned}
B_1 &= \bigl\{\omega\colon\ 0\leqslant \min\bigl(Y_1(\omega), Y_2(\omega)\bigr)\vee \max\bigl(Y_1(\omega), Y_2(\omega)\bigr)\leqslant 0\bigr\}, \\
B_2 &= \bigl\{\omega\colon\ \bigl(Y_1(\omega) < 0 < Y_2(\omega)\bigr)\vee\bigl(Y_2(\omega) < 0 < Y_1(\omega)\bigr)\bigr\}.
\end{aligned}$$

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

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

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

$$\begin{aligned}&\int_{B_1} \max\bigl(Y_1(\omega) + Y_2(\omega), 0\bigr) P(\mathrm{d}\omega) \\ &\quad = \int_{B_1} \max\bigl(Y_1(\omega), 0\bigr) P(\mathrm{d}\omega) + \int_{B_1} \max\bigl(Y_2(\omega), 0\bigr) P(\mathrm{d}\omega).\end{aligned} \tag{F.1}$$

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

$$\begin{aligned}&\int_{B_2} \max\bigl(Y_1(\omega) + Y_2(\omega), 0\bigr) P(\mathrm{d}\omega) \\ &\quad \leqslant \int_{B_2} \max\bigl(Y_1(\omega), 0\bigr) P(\mathrm{d}\omega) + \int_{B_2} \max\bigl(Y_2(\omega), 0\bigr) P(\mathrm{d}\omega).\end{aligned} \tag{F.2}$$

Summing left and right hand sides of Eqs. (F.1) and (F.2), we obtain

$$\mathrm{E}\bigl[(Y_1 + Y_2)^+\bigr] \leqslant \mathrm{E}\bigl[Y_1^+\bigr] + \mathrm{E}\bigl[Y_2^+\bigr]. \tag{F.3}$$

If $P(B_2) > 0$, then the inequality in Eq. (F.2) is strict, and therefore the inequality in Eq. (F.3) is strict as well.

Appendix G. 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 Eq. (17). As $n \to \infty$, L_n converges to $\mu(X)$, so asymptotic EEL is equal to the value of c solving

$$\theta = \mathrm{E}\bigl[\bigl(\mu(X) - c\bigr)^+\bigr] = \int_{\mu^{-1}(c)}^{\infty} \bigl(\mu(x) - c\bigr) h(x)\,\mathrm{d}x, \tag{G.1}$$

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

$$\theta = (EL - c)\bigl(1 - H\bigl(\mu^{-1}(c)\bigr)\bigr) + \frac{EL \cdot w}{\Gamma(1 + 1/\sigma^2)}\bigl(\mu^{-1}(c)/\sigma^2\bigr)^{1/\sigma^2} \exp\left(\frac{-\mu^{-1}(c)}{\sigma^2}\right), \tag{G.2}$$

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

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

The gamma cdf is available in nearly all numerical packages. 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 Eq. (G.2) simplifies to

$$\theta = w \cdot EL \cdot \exp\bigl(-\mu^{-1}(c)\bigr).$$

This yields the closed-form solution

$$\mathrm{EEL}_\theta[L_\infty] = c = EL - w \cdot EL \cdot \bigl(1 + \ln(\theta) - \ln(w \cdot EL)\bigr).$$

References

Abramowitz, M., Stegun, I.A., 1968. Handbook of Mathematical Functions. National Bureau of Standards.

Acerbi, C., Tasche, D., 2002. On the coherence of expected shortfall. J. Banking Finance 26, 1487–1503.

Artzner, P., Delbaen, F., Eber, J.-M., Heath, D., 1999. Coherent measures of risk. Math. Finance 9, 203–228.

Basel Committee on Bank Supervision, 1999. Credit risk modelling: Current practices and applications. Technical report. Bank for International Settlements.

Basel Committee on Bank Supervision, 2001. The Internal Ratings-Based approach: Supporting document to the New Basel Capital Accord. Technical report. Bank for International Settlements.

Billingsley, P., 1995. Probability and Measure, 3rd edition. Wiley, New York.

Bürgisser, P., Kurth, A., Wagner, A., 2001. Incorporating severity variations into credit risk. Risk 3, 5–31.

Carey, M., 2001. Dimensions of credit risk and their relationship to economic capital requirements. In: Mishkin, F.S. (Ed.), Prudential Supervision: What Works and What Doesn't. Univ. of Chicago Press.

Credit Suisse Financial Products, 1977. CreditRisk$^+$: A Credit Risk Management Framework. Credit Suisse Financial Products.

Finger, C.C., 1999. Conditional approaches for CreditMetrics portfolio distributions. CreditMetrics Monitor, 14–33.

Frey, R., McNeil, A.J., 2002. VaR and expected shortfall in portfolios of dependent credit risks: conceptual and practical insights. J. Banking Finance 26, 1317–1334.

Frye, J., 2000. Collateral damage: A source of systematic credit risk. Risk 13.

Gordy, M.B., 2000. A comparative anatomy of credit risk models. J. Banking Finance 24, 119–149.

Gordy, M., Jones, D., 2003. Random tranches. Risk 16, 78–83.

Gouriéroux, C., Laurent, J., Scaillet, O., 2000. Sensitivity analysis of values at risk. J. Empirical Finance 7, 225–245.

Gupton, G.M., Finger, C.C., Bhatia, M., 1997. CreditMetrics—Technical Document. J.P. Morgan, New York.

Hill, G., Davis, A., 1968. Generalized asymptotic expansions of Cornish–Fisher type. Ann. Math. Statist. 39, 1264–1273.

Jones, D., 2000. Emerging problems with the Basel capital accord: regulatory capital arbitrage and related issues. J. Banking Finance 24, 35–58.

Knopp, K., 1956. Infinite Sequences and Series. Dover, New York.
Martin, R., Wilde, T., 2002. Unsystematic credit risk. Risk 15, 123–128.
Petrov, V.V., 1995. Limit Theorems of Probability Theory. Oxford Univ. Press.
Pykhtin, M., Dev, A., 2002. Credit risk in asset securitizations: analytical model. Risk 15, S16–S20.
Pykhtin, M., Dev, A., 2003. Coarse-grained CDOs. Risk 16, 113–116.
Rockafellar, R.T., Uryasev, S., 2002. Conditional value-at-risk for general loss distributions. J. Banking Finance 26, 1443–1471.
Szegö, G.P., 2002. Measures of risk. J. Banking Finance 26, 1253–1272.
Tasche, D., 2000. Conditional expectation as quantile derivative. Working paper. TU München.
Tasche, D., 2001. Calculating value-at-risk contributions in CreditRisk$^+$. Working paper. ETH Zentrum.
Tasche, D., 2002. Expected shortfall and beyond. J. Banking Finance 26, 1519–1533.
Treacy, W.F., Carey, M.S., 1998. Credit risk rating at large US banks. Fed. Reserve Bull. 84, 897–921.
Vasicek, O.A., 1997. The loan loss distribution. Technical report. KMV Corporation.
Wilde, T., 2001. Probing granularity. Risk 14, 103–106.