A Comparative Anatomy of Credit Risk Models

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

December 8, 1998

Abstract

Within the past two years, important advances have been made in modeling credit risk at the portfolio level. Practitioners and policy makers have invested in implementing and exploring a variety of new models individually. Less progress has been made, however, with comparative analyses. Direct comparison often is not straightforward, because the different models may be presented within rather different mathematical frameworks.

This paper offers a comparative anatomy of two especially influential benchmarks for credit risk models, J.P. Morgan's CreditMetrics and Credit Suisse Financial Product's CreditRisk$^+$. We show that, despite differences on the surface, the underlying mathematical structures are similar. The structural parallels provide intuition for the relationship between the two models and allow us to describe quite precisely where the models differ in functional form, distributional assumptions, and reliance on approximation formulae. We then design simulation exercises which evaluate the effect of each of these differences individually.

JEL Codes: G31, C15, G11

Over the past decade, financial institutions have developed and implemented a variety of sophisticated models of value-at-risk for market risk in trading portfolios. These models have gained acceptance not only among senior bank managers, but also in amendments to the international bank regulatory framework. Much more recently, important advances have been made in modeling credit risk in lending portfolios. The new models are designed to quantify credit risk on a portfolio basis, and thus have application in control of risk concentration, evaluation of return on capital at the customer level, and more active management of credit portfolios. Future generations of today's models may one day become the foundation for measurement of regulatory capital adequacy.

Two of the models, J.P. Morgan's CreditMetrics and Credit Suisse Financial Product's CreditRisk$^+$, have been released freely to the public since 1997 and have quickly become influential benchmarks. Practitioners and policy makers have invested in implementing and exploring each of the models individually, but have made less progress with comparative analyses. The two models are intended to measure the same risks, but impose different restrictions and distributional assumptions, and suggest different techniques for calibration and solution. Thus, given the same portfolio of credit exposures, the two models will, in general, yield differing evaluations of credit risk. Determining which features of the models account for differences in output would allow us a better understanding of the sensitivity of the models to the particular assumptions they employ.

Unfortunately, direct comparison of the models is not straightforward, because the two models are presented within rather different mathematical frameworks. The CreditMetrics model is familiar to econometricians as an ordered probit model. Credit events are driven by movements in underlying unobserved latent variables. The latent variables are assumed to depend on external "risk factors." Common dependence on the same risk factors gives rise to correlations in credit events across obligors. The CreditRisk$^+$ model is based instead on insurance industry models of event risk. Instead of a latent variable, each obligor has a default probability. The default probabilities are not constant over time, but rather increase or decrease in response to background macroeconomic factors. To the extent that two obligors are sensitive to the same set of background factors, their default probabilities will move together. These co-movements in probability give rise to correlations in defaults. CreditMetrics and CreditRisk$^+$ may serve essentially the same function, but they appear to be constructed quite differently.

This paper offers a comparative anatomy of CreditMetrics and CreditRisk$^+$. We show that, despite differences on the surface, the underlying mathematical structures are similar. The structural parallels provide intuition for the relationship between the two models and allow us to describe quite precisely where the models differ in functional form, distributional assumptions, and reliance on approximation formulae. We can then design simulation exercises which evaluate the effect of these differences individually.

We proceed as follows. Section 1 presents a summary of the CreditRisk$^+$ model, and introduces a restricted version of CreditMetrics. The restrictions are imposed to facilitate direct comparison of CreditMetrics and CreditRisk$^+$. While some of the richness of the full CreditMetrics implementation is sacrificed, the essential mathematical characteristics of the model are preserved. Our

1

method of comparative anatomy is developed in Section 2. We show how the restricted version of CreditMetrics can be run through the mathematical machinery of CreditRisk$^+$, and vice versa.

Comparative simulations are developed in Section 3. Care is taken to construct portfolios with quality and loan size distributions similar to real bank portfolios, and to calibrate correlation parameters in the two models in an empirically plausible and mutually consistent manner. The robustness of the conclusions of Section 3 to our methods of portfolio construction and parameter calibration are explored in Section 4. An especially striking result from the simulations is the sensitivity of CreditRisk$^+$ results to the shape of the distribution of a systematic risk factor. The causes for and consequences of this sensitivity are explored in Section 5. We conclude with a summary of the main results of the simulations.

1   Summary of the models

This section offers an introduction to CreditRisk$^+$ and CreditMetrics. The discussion of CreditRisk$^+$ merely summarizes the derivation presented in Credit Suisse Financial Products (1997, Appendix A). Our presentation of CreditMetrics sets forth a restricted version of the full model described in the CreditMetrics Technical Document (Gupton, Finger and Bhatia 1997). Our choice of notation is intended to facilitate comparison of the models, and may differ considerably from what is used in the original manuals.

1.1   Summary of CreditRisk$^+$

CreditRisk$^+$ is a model of *default risk*. Each obligor has only two possible end-of-period states, default and non-default. In the event of default, the lender suffers a loss of fixed size; this is the lender's *exposure* to the obligor. The distributional assumptions and functional forms imposed by CreditRisk$^+$ allow the distribution of total portfolio losses to be calculated in a convenient analytic form.

Default correlations in CreditRisk$^+$ are assumed to be driven entirely by a vector of K "risk factors" $x = (x_1, \ldots, x_K)$. Conditional on x, defaults of individual obligors are assumed to be independently distributed Bernoulli draws. The conditional probability $p_i(x)$ of drawing a default for obligor i is a function of the rating grade $\zeta(i)$ of obligor i, the realization of risk factors x, and the vector of "factor loadings" $(w_{i1}, \ldots, w_{iK})$ which measure the sensitivity of obligor i to each of the risk factors. CreditRisk$^+$ specifies this function as

$$p_i(x) = \bar{p}_{\zeta(i)} \left(\sum_{k=1}^{K} x_k w_{ik} \right) \tag{1}$$

where $\bar{p}_\zeta$ is the unconditional default probability for a grade ζ obligor, and the x are positive-valued with mean one. The intuition behind this specification is that the risk factors x serve to "scale up" or "scale down" the unconditional $\bar{p}_\zeta$. A high draw of x_k (over one) increases the probability of

2

default for each obligor in proportion to the obligor's weight w_{ik} on that risk factor; a low draw of x_k (under one) scales down all default probabilities. The weights w_{ik} are required to sum to one for each obligor, which guarantees that $\mathrm{E}[p_i(x)] = \bar{p}_{\zeta(i)}$.

Rather than calculating the distribution of defaults directly, CreditRisk$^+$ calculates the *probability generating function* ("pgf") for defaults. The pgf $\mathcal{F}_\kappa(z)$ of a discrete random variable κ is a function of an auxilliary variable z such that the probability that $\kappa = n$ is given by the coefficient on z^n in the polynomial expansion of $\mathcal{F}_\kappa(z)$. The pgf, which is essentially a discrete random variable analog to the moment generating function, has two especially useful properties:[1]

- If κ_1 and κ_2 are independent random variables, then the pgf of the sum $\kappa_1 + \kappa_2$ is equal to the product of the two pgfs.

- If $\mathcal{F}_\kappa(z|x)$ is the pgf of κ conditional on x, and x has distribution function $H(x)$, then the unconditional pgf is simply $\mathcal{F}_\kappa(z) = \int_x \mathcal{F}_\kappa(z|x) dH(x)$.

We first derive the conditional pgf $\mathcal{F}(z|x)$ for the total number of defaults in the portfolio given realization x of the risk factors. For a single obligor i, this is the Bernoulli($p_i(x)$) pgf:

$$\mathcal{F}_i(z|x) = (1 - p_i(x) + p_i(x)z) = (1 + p_i(x)(z - 1)).$$

Using the approximation formula $\log(1 + y) \approx y$ for $y \approx 0$, we can write

$$\mathcal{F}_i(z|x) = \exp(\log(1 + p_i(x)(z - 1))) \approx \exp(p_i(x)(z - 1)). \tag{2}$$

We refer to this step as the "Poisson approximation" because the expression on the right hand side is the pgf for a random variable distributed Poisson($p_i(x)$). The intuition is that, as long as $p_i(x)$ is small, we can ignore the constraint that a single obligor can default only once, and represent its default event as a Poisson random variable rather than as a Bernoulli. The exponential form of the Poisson pgf is essential to the computational facility of the model.

Conditional on x, default events are independent across obligors, so the pgf of the sum of obligor defaults is the product of the individual pgfs:

$$\mathcal{F}(z|x) = \prod_i \mathcal{F}_i(z|x) \approx \prod_i \exp(p_i(x)(z - 1)) = \exp(\mu(x)(z - 1)) \tag{3}$$

where $\mu(x) \equiv \sum_i p_i(x)$.

To get the unconditional probability generating function $\mathcal{F}(z)$, we integrate out the x. The risk factors in CreditRisk$^+$ are assumed to be independent gamma-distributed random variables with mean one and variance σ_k^2, $k = 1, \ldots, K$.[2] See Appendix A on the properties and parameterization

---

[1] See Johnson and Kotz (1969, §2.2) for further discussion on probability generating functions.

[2] This is a variant on the presentation in the CreditRisk$^+$ manual, in which x_k has mean μ_k and variance σ_k^2, and the conditional probabilities are given by $p_i(x) = \bar{p}_{\zeta(i)}(\sum w_{ik}(x_k/\mu_k))$. In our presentation, the constants $1/\mu_k$ are absorbed into the normalized x_k without any loss of generality.

of the gamma distribution. It is straightforward to show that

$$\mathcal{F}(z) = \prod_{k=1}^{K} \left(\frac{1 - \delta_k}{1 - \delta_k z} \right)^{1/\sigma_k^2} \quad \text{where } \delta_k \equiv \frac{\sigma_k^2 \mu_k}{1 + \sigma_k^2 \mu_k} \text{ and } \mu_k \equiv \sum_i w_{ik} \bar{p}_{\zeta(i)}. \tag{4}$$

The form of this pgf shows that the total number of defaults in the portfolio is a sum of K independent negative binomial variables.

The final step in CreditRisk$^+$ is to obtain the probability generating function $\mathcal{G}(z)$ for losses. Assume loss given default is a constant fraction λ of loan size. Let L_i denote the loan size for obligor i. In order to retain the computational advantages of the discrete model, we need to express the loss exposure amounts λL_i as integer multiples of a fixed unit of loss (e.g., one million dollars). The base unit of loss is denoted ν_0 and its integer multiples are called "standardized exposure" levels. The standardized exposure for obligor i, denoted $\nu(i)$, is equal to $\lambda L_i / \nu_0$ rounded to the nearest integer.

Let $\mathcal{G}_i$ denote the probability generating function for losses on obligor i. The probability of a loss of $\nu(i)$ units on a portfolio consisting only of obligor i must equal the probability that i defaults, so $\mathcal{G}_i(z|x) = \mathcal{F}_i(z^{\nu(i)}|x)$. We use the conditional independence of the defaults to obtain the conditional pgf for losses in the entire portfolio as

$$\mathcal{G}(z|x) = \prod_i \mathcal{G}_i(z|x) = \exp \left(\sum_{k=1}^{K} x_k \sum_i \bar{p}_{\zeta(i)} w_{ik}(z^{\nu(i)} - 1) \right).$$

As before, we integrate out the x and rearrange to arrive at

$$\mathcal{G}(z) = \prod_{k=1}^{K} \left(\frac{1 - \delta_k}{1 - \delta_k \mathcal{P}_k(z)} \right)^{1/\sigma_k^2} \quad \text{where} \quad \mathcal{P}_k(z) \equiv \frac{1}{\mu_k} \sum_i w_{ik} \bar{p}_{\zeta(i)} z^{\nu(i)} \tag{5}$$

and δ_k and μ_k are as defined in equation (4).

The unconditional probability that there will be n units of ν_0 loss in the total portfolio is given by the coefficient on z^n in the Taylor series expansion of $\mathcal{G}(z)$. The CreditRisk$^+$ manual (§A.10) provides the recurrence relation used to calculate these coefficients.

1.2 A restricted version of CreditMetrics

The CreditMetrics model for credit events is familiar to economists as an ordered probit. Associated with obligor i is an unobserved latent random variable y_i. The state of obligor i at the risk-horizon depends on the location of y_i relative to a set of "cut-off" values. In the full version of the model, the cut-offs divide the real number line into "bins" for each end-of-period rating grade. CreditMetrics thereby captures not only defaults, but migrations across non-default grades as well. Given a set of forward credit spreads for each grade, CreditMetrics can then estimate a distribution over the change in mark-to-market value attributable to portfolio credit risk.

In this section, we present a restricted version of CreditMetrics. To allow more direct comparison

4

with CreditRisk$^+$, we restrict the set of outcomes to two states: default and non-default. In the event of default, we assume loss is a fixed fraction λ of the face value. This represents a second significant simplification of the full CreditMetrics implementation, which allows idiosyncratic risk in recoveries.[3] In the non-default state, the loan retains its book value. Thus, our restricted version of CreditMetrics is a model of book value losses, rather than of changes in market value. In the discussion below, the restricted CreditMetrics will be designated as "CM2S" ("CreditMetrics two-state") whenever distinction from the full CreditMetrics model needs emphasis.

The latent variables y_i are taken to be linear functions of risk factors x and idiosyncratic effects ϵ_i:

$$y_i = xw_i + \eta_i \epsilon_i. \tag{6}$$

The vector of factor loadings w_i determines the relative sensitivity of obligor i to the risk factors, and the weight η_i determines the relative importance of idiosyncratic risk for the obligor. The x are assumed to be normally distributed with mean zero and variance-covariance matrix Ω.[4] Without loss of generality, assume there are ones on the diagonal of Ω, so the marginal distributions are all $N(0,1)$. The ϵ_i are assumed to be iid $N(0,1)$. Again without loss of generality, it is imposed that y_i has variance 1 (i.e., that $w_i' \Omega w_i + \eta_i^2 = 1$). Associated with each start-of-period rating grade ζ is a cut-off value C_ζ. When the latent variable y_i falls under the cut-off $C_{\zeta(i)}$, the obligor defaults. That is, default occurs if

$$xw_i + \eta_i \epsilon_i < C_{\zeta(i)}. \tag{7}$$

The C_ζ values are set so that the unconditional default probability for a grade ζ obligor is $\bar{p}_\zeta$, i.e., so that $\bar{p}_\zeta = \Phi(C_\zeta)$, where Φ is the standard normal cdf and the $\bar{p}_\zeta$ are defined as in Section 1.1.

The model is estimated by Monte Carlo simulation. To obtain a single "draw" for the portfolio, we first draw a single vector x as a multivariate $N(0, \Omega)$ and a set of iid $N(0,1)$ idiosyncratic ϵ. We form the latent y_i for each obligor, which are compared against the cut-off values $C_{\zeta(i)}$ to determine default status D_i (one for default, zero otherwise). Portfolio loss for this draw is given by $\sum_i D_i \lambda L_i$. To estimate a distribution of portfolio outcomes, we repeat this process many times. The portfolio losses obtained in each draw are sorted to form a cumulative distribution for loss. For example, if the portfolio is simulated 100,000 times, then the estimated 99.5th percentile of the loss distribution is given by the 99500th element of the sorted loss outcomes.

---

[3]The full CreditMetrics also accommodates more complex asset types, including loan commitments and derivatives contracts. See the CreditMetrics Technical Document, Chapter 4, and Finger (1998). These features are not addressed in this paper.

[4]In the CreditMetrics Technical Document, it is recommended that the x be taken to be stock market indexes, because the ready availability of historical data on stock indexes simplifies calibration of the covariance matrix Ω and the weights w_i. The mathematical framework of the model, however, imposes no specific identity on the x.

2 Mapping between the models

Presentation of the restricted version of CreditMetrics and the use of a similar notation in outlining the models both serve to emphasize the fundamental similarities between CreditMetrics and CreditRisk$^+$. Nonetheless, there remain substantial differences in the mathetical methods used in each, which tend to obscure comparison of the models. In the first two parts of this section, we map each model into the mathematical framework of the other. We conclude this section with a comparative analysis in which fundamental differences between the models in functional form and distributional assumptions are distinguished from differences in technique for calibration and solution.

2.1 Mapping CreditMetrics to the CreditRisk$^+$ framework

To map the restricted CM2S model into the CreditRisk$^+$ framework, we need to derive the implied conditional default probability function $p_i(x)$ used in equation (2). Conditional on x, rearrangement of equation (7) shows that obligor i defaults if and only if $\epsilon_i < ((C_{\zeta(i)} - xw_i)/\eta_i$. Because the ϵ_i are standard normal variates, default occurs with conditional probability

$$p_i(x) = \Phi((C_{\zeta(i)} - xw_i)/\eta_i). \tag{8}$$

The CreditRisk$^+$ methodology can now be applied in a straightforward manner. Conditional on x, default events are independent across obligors. Therefore, the conditional probability generating function for defaults, $\mathcal{F}(z|x)$, takes on exactly the same Poisson approximation form as in equation (3). To get the unconditional probability generating function $\mathcal{F}(z)$, we integrate out the x:

$$\mathcal{F}(z) = \int_{-\infty}^{\infty} \mathcal{F}(z|x)\phi_{\Omega}(x)dx$$

where ϕ_{Ω} is the multivariate $N(0, \Omega)$ pdf. The unconditional probability that exactly n defaults will occur in the portfolio is given by the coefficient on z^n in the Taylor series expansion of $\mathcal{F}(z)$:

$$\begin{aligned}
\mathcal{F}(z) &= \int_{-\infty}^{\infty} \sum_{n=0}^{\infty} \exp(-\mu(x))\frac{\mu(x)^n z^n}{n!}\phi_{\Omega}(x)dx \\
&= \sum_{n=0}^{\infty} \frac{1}{n!}\left(\int_{-\infty}^{\infty} \exp(-\mu(x))\mu(x)^n \phi_{\Omega}(x)dx\right) z^n
\end{aligned} \tag{9}$$

where $\mu(x) \equiv \sum_i p_i(x)$. These K-dimensional integrals are analytically intractable (even when the number of systematic factors $K = 1$), and in practice would be solved using Monte Carlo techniques.[5]

The remaining steps in CreditRisk$^+$ would follow similarly. That is, one would round loss

---

[5]Note that the integrals differ only in n, so a single set of Monte Carlo draws for x allows successive solution of the integrals via a simple recurrence relation. This technique is fast relative to standard Gaussian quadrature.

exposures to integer multiples of a base unit ν_0, apply the rule $\mathcal{G}_i(z|x) = \mathcal{F}_i(z^{\nu(i)}|x)$ and multiply pgfs across obligors to get $\mathcal{G}(z|x)$. To get the unconditional pgf of losses, we integrate out the x as in equation (9). The result would be computationally unwieldy, but application of the method is conceptually straightforward.

2.2 Mapping CreditRisk$^+$ to the CreditMetrics framework

Translating in the opposite direction is equally straightforward. To go from CreditRisk$^+$ into the CM2S framework, we assign to obligor i a latent variable y_i defined by:

$$y_i = \left(\sum_{k=1}^{K} x_k w_{ik} \right)^{-1} \epsilon_i. \tag{10}$$

The x_k and w_{ik} are the same gamma-distributed risk factors and factor loadings used in CreditRisk$^+$. The idiosyncratic risk factors ϵ_i are independently and identically distributed Exponential with parameter 1. Obligor i defaults if and only if $y_i < \bar{p}_{\zeta(i)}$. Observe that the conditional probability of default is given by:

$$
\begin{aligned}
\Pr(y_i < \bar{p}_{\zeta(i)}|x) &= \Pr\left(\epsilon_i < \bar{p}_{\zeta(i)} \sum_{k=1}^{K} x_k w_{ik}|x \right) \\
&= 1 - \exp\left(-\bar{p}_{\zeta(i)} \sum_{k=1}^{K} x_k w_{ik} \right) \\
&\approx \bar{p}_{\zeta(i)} \sum_{k=1}^{K} x_k w_{ik} = p_i(x)
\end{aligned}
\tag{11}
$$

where the second line follows using the cdf for the exponential distribution, and the last line relies on the same approximation formula as equation (2). The unconditional probability of default is simply $\bar{p}_{\zeta(i)}$, as required.

In the ordinary CreditMetrics specification, the latent variable is a linear sum of normal random variables. When CreditRisk$^+$ is mapped to the CreditMetrics framework, the latent variable takes a multiplicative form, but the idea is the same.[6] In CreditMetrics, the cut-off values C_ζ are determined as functions of the associated unconditional default probabilities $\bar{p}_\zeta$. Here, the cut-off values are simply the $\bar{p}_\zeta$. Other than these differences in form, the process is identical. A single portfolio simulation would consist of a single random draw of sector risk factors and a single vector of random draws of idiosyncratic risk factors. From these, the obligors' latent variables are calculated, and these in turn determine default events.[7]

---

[6]One could quasi-linearize equation (10) by taking logs, but little would be gained. The log of an exponential random variable does not itself have a well-known distributional form, and the log of the weighted sum of x variables would not simplify.

[7]In practice, it is faster and more accurate to simulate outcomes as independent Bernoulli($p_i(x)$) draws, which is the method used in Section 5. The latent variable method is presented only to emphasize the structural similarities between the models.

2.3 Essential and inessential differences between the models

CreditMetrics and CreditRisk$^+$ differ in distributional assumptions and functional forms, solution techniques, suggested methods for calibration, and mathematical language. As the preceding analysis makes clear, only the differences in distributional assumptions and functional forms are fundamental. Each model can be mapped into the mathematical language of the other, which demonstrates that the difference between the latent variable representation of CreditMetrics and the covarying default probabilities of CreditRisk$^+$ is one of presentation and not substance. Similarly, methods for parameter calibration suggested in model technical documents are helpful to users, but not in any way intrinsic to the models.

By contrast, distributional assumptions and functional forms are model primitives. In each model, the choice of distribution for the systemic risk factors x and the functional form for the conditional default probabilities $p_i(x)$ together give shape to the joint distribution over obligor defaults in the portfolio. The CreditMetrics specifications of normally distributed x and of equation (8) for the $p_i(x)$ may be somewhat arbitrary, but nonetheless strongly influence the results.[8] One could substitute, say, a multivariate t distribution for the normal distribution, and still employ the Monte Carlo methodology of CreditMetrics. However, even if parameters were recalibrated to yield the same mean and variance of portfolio loss, the overall shape of the loss distribution would differ, and therefore the tail percentile values would change as well. The choice of the gamma distribution and the function form for conditional default probabilities given by equation (1) are similarly characteristic of CreditRisk$^+$. Indeed, in Section 5 we will show how small deviations from the gamma specification lead to significant differences in tail percentile values in a generalized CreditRisk$^+$ framework.

Remaining differences between the two models are attributable to differences in solution method. The Monte Carlo method of CreditMetrics is flexible but computationally intensive. CreditRisk$^+$ offers the efficiency of a closed-form solution, but at the expense of additional restrictions or approximations. In particular,

- CreditMetrics allows naturally for multi-state outcomes and for uncertainty in recoveries, whereas the closed-form CreditRisk$^+$ is a two-state model with fixed recovery rates.[9]

- CreditRisk$^+$ imposes a "Poisson approximation" on the conditional distribution of defaults.

- CreditRisk$^+$ rounds each obligor's loan loss exposure to the nearest element in a finite set of

---

[8]Recall that equation (8) follows from the specification in equation (6) of latent variable y_i and the assumption of normally distribution idiosyncratic ϵ_i. Therefore, this discussion incorporates those elements of the CreditMetrics model.

[9]In principle, both restrictions in CreditRisk$^+$ can be relaxed without resorting to a Monte Carlo methodology. It is feasible to introduce idiosyncratic risk in recoveries to CreditRisk$^+$, but probably at the expense of the computational facility of the model. CreditRisk$^+$ also can be extended to, say, a three-state model in which the third state represents a severe downgrade (short of default). However, CreditRisk$^+$ cannot capture the exclusive nature of the outcomes (i.e., one cannot impose the mutual exclusivity of a severe downgrade and default). For this reason, as well as the Poisson approximation, the third state would need to represent a low probability event. Thus, it would be impractical to model ordinary rating migrations in CreditRisk$^+$.

values.

Using the techniques of Section 2.2, it is straightforward to construct a Monte Carlo version of CreditRisk$^+$ which avoids Poisson and loss exposure approximations and allows recovery risk. It is less straightforward but certainly possible to create a Monte Carlo multi-state generalization of CreditRisk$^+$. Because computational convenience may be a significant advantage of CreditRisk$^+$ for some users, the effect of the Poisson and loss exposure approximations on the accuracy of CreditRisk$^+$ results will be examined in Sections 4 and 5. However, the effect of multi-state outcomes and recovery uncertainty on the distribution of credit loss will be left for future study. Therefore, our simulations will study only the restricted CM2S version of CreditMetrics.

It is worth noting that there is no real loss of generality in the assumption of independence across sector risk factors in CreditRisk$^+$. In each model, the vector of factor loadings (w) is free, up to a scaling restriction. In CreditMetrics, the sector risk factors x could be orthogonalized and the correlations incorporated into the w.[10] However, the need to impose orthogonality in CreditRisk$^+$ does imply that greater care must be given to identifying and calibrating sectoral risks in that model.

3 Calibration and main simulation results

The remaining sections of this paper study the two models using comparative simulations. The primary goal is to develop reliable intuition for how the two models will differ when applied to real world portfolios. In pursuit of this goal, we also will determine the parameters or portfolio characteristics to which each model is most sensitive. Emphasis is placed on relevance and robustness. By relevance, we mean that the simulated portfolios and calibrated parameters ought to resemble their real world counterparts closely enough for conclusions to be transferable. By robustness we mean that the conclusions ought to be qualitatively valid over an empirically relevant range of portfolios.

This section will present our main simulations. First, in Section 3.1, we construct a set of "test deck" portfolios. All assets are assumed to be ordinary term loans. The size distribution of loans and their distribution across S&P rating grades are calibrated using data from two large samples of mid-sized and large corporate loans. Second, in Section 3.2, default probabilities and correlation structures in each model are calibrated using historical default data from the S&P ratings universe. We calibrate each model to a one year risk-horizon. The main simulation results are presented in Section 3.3.

3.1 Portfolio construction

Construction of our simulated loan portfolios requires choices along three dimensions. The first is *credit quality*, i.e., the portion of total dollar outstandings in each rating grade. The second is *obligor count*, i.e., the total number of obligors in the portfolio. The third is *concentration*, i.e.,

---

[10]In this case, the original weights w would be replaced by $\Omega^{1/2}w$.

the distribution of dollar outstandings within a rating grade across the obligors within that grade. Note that the total portfolio dollar outstandings is immaterial, because losses will be calculated as a percentage of total outstandings.

The range of plausible credit quality is represented by four credit quality distributions, which are labelled "High," "Average," "Low," and "Very Low." The first three distributions are constructed using data from internal Federal Reserve Board surveys of large banking organizations.[11] The "Average" distribution is the average distribution across the surveyed banks of total outstandings in each S&P grade. The "High" and "Low" distributions are drawn from the higher and lower quality distributions found among the banks in the sample. The "Very Low" distribution is not found in the Federal Reserve sample, but is intended to represent a very weak large bank loan portfolio during a recession. Speculative grade (BB and below) loans account for half of outstandings in the "Average" portfolio, and 25%, 78% and 83% in the "High," "Low," and "Very Low" quality portfolios, respectively. The distributions are depicted in Figure 1.

Realistic calibration of obligor count is likely to depend not only on the size of the hypothetical bank, but also on the bank's business focus. A very large bank with a strong middle-market business might have tens of thousands of rated obligors in its commercial portfolio. A bank of the same size specializing in the large corporate market might have only a few thousand. For the "base case" calibration, we set $N = 5000$. To establish robustness of the conclusions to the choice of N, we model portfolios of 1000 and 10,000 obligors as well. In all simulations, we assume each obligor is associated with only one loan in the portfolio.

Portfolio concentration is calibrated in two stages. First, we divide the total number of obligors N across the rating grades. Second, for each rating grade, we determine how the total exposure within the grade is distributed across the number of obligors in the grade. In both stages, distributions are calibrated using the Society of Actuaries (1996) (hereafter cited as "SoA") sample of mid-sized and large private placement loans (see also Carey 1998).

Let q_ζ be the dollar volume of exposure to rating grade ζ as a share of total portfolio exposure (determined by the chosen "credit quality distribution"), and let v_ζ be the mean book value of loans in rating grade ζ in the SoA data. We determine n_ζ, the number of obligors in rating grade ζ, by imposing

$$q_\zeta = n_\zeta v_\zeta / \sum_g n_g v_g \tag{12}$$

for all ζ. That is, the n_ζ are chosen so that, in a portfolio with mean loan sizes in each grade matching the SoA mean sizes, the exposure share of that grade matches the desired share q_ζ. The equations of form (12) are easily transformed into a set of six linearly independent equations and seven unknown n_ζ values (using the S&P eight grade scale). Given the restriction $\sum n_\zeta = N$, the vector n is uniquely determined.[12] Table 1 shows the values of the vector n associated with each

---

[11] Each bank provided the amount outstanding, by rating grade, in its commericial loan book.

[12] After solving the system of linear equations, the values of n_ζ are rounded to whole numbers.

Figure 1: Credit Quality Distributions
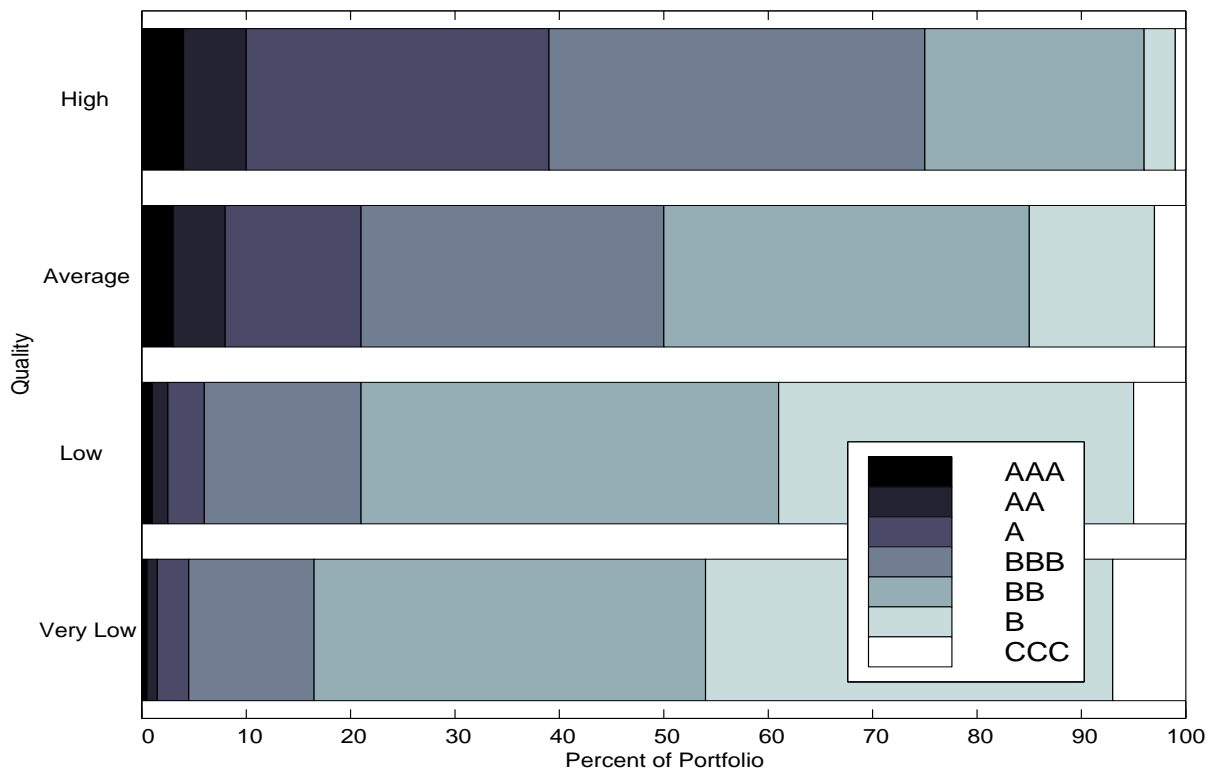
credit quality distribution when $N = 5000$.

Table 1: Number of Obligors in Each Rating Grade

| | Portfolio Credit Quality | | | |
	High	Average	Low	Very Low
AAA	191	146	50	25
AA	295	250	77	51
A	1463	669	185	158
BBB	1896	1558	827	660
BB	954	1622	1903	1780
B	136	556	1618	1851
CCC	65	199	340	475
Total	5000	5000	5000	5000

The final step is to distribute the q_ζ share of total loan exposure across the n_ζ obligors in each grade. For our "base case" calibration, the distribution within grade ζ is chosen to match (up to a scaling factor) the distribution for grade ζ exposures in the SoA data. The SoA loans in grade ζ are sorted from smallest to largest, and used to form a cumulative distribution H_ζ. The size of the j^{th} exposure, $j = 1, \ldots, n_\zeta$, is set to the $(j - 1/2)/n_\zeta$ percentile of H_ζ.[13] Finally, the n_ζ loan sizes are normalized to sum to q_ζ. This method ensures that the shape of the distribution of loan sizes will not be sensitive to the choice of n_ζ (unless n_ζ is very small). As an alternative to the base case, we also model a portfolio in which all loans within a rating grade are equal-sized, i.e., each loan in grade ζ is of size q_ζ/n_ζ.

In all simulations below, it is assumed that loss given default is a fixed proportion $\lambda = 0.3$ of book value, which is consistent with historical loss given default experience for senior unsecured bank loans.[14] Percentile values on the simulated loss distributions are directly proportional to λ. Holding fixed all other model parameters and the permitted probability of bank insolvency, the required capital given a loss rate of, say, $\lambda = 0.45$ would simply be 1.5 times the required capital given $\lambda = 0.3$.

It should be emphasized that the SoA data are used to impose *shape*, but not *scale*, on the distributions of loan sizes. The ratio of the mean loan size in grade ζ_1 to the mean loan size in grade ζ_2 is determined by the corresponding ratio in the SoA sample. Within each grade, SoA data determine the ratio of any two percentile values of loan size (e.g., the 75th percentile to the median). However, measures of portfolio concentration (e.g., the ratio of the sum of the largest j loans to the total portfolio value) depend strongly on the choice of N, and thus not only on SoA

---

[13]For example, say $n_\zeta = 200$ and there are 6143 loans in the grade ζ SoA sample. To set the size of, say, the seventh ($j = 7$) simulated loan, we calculated the index $(j - 1/2) \cdot 6143/n_\zeta \approx 199.65$. The loan size is then formed as an interpolated value between the 199^{th} and 200^{th} loans in the SoA vector.

[14]See the CreditMetrics Technical Document, §7.1.2.

sample.

Finally, CreditRisk$^+$ requires a discretization of the distribution of exposures, i.e., the selection of the base unit of loss ν_0. In the main set of simulations, we will set ν_0 to λ times the fifth percentile value of the distribution of loan sizes. In Section 4, it will be shown that simulation results are quite robust to the choice of ν_0.

3.2 Default probabilites and correlations

In any model of portfolio credit risk, the structure of default rate correlations is an important determinant of the distribution of losses. Special attention must therefore be given to mutually consistent calibration of parameters which determine default correlations. In the exercises below, we calibrate CreditMetrics and CreditRisk$^+$ to yield the same unconditional expected default rate for an obligor of a given rating grade, and the same default correlation between any two obligors within a single rating grade.

For simplicity, we assume a single systematic risk factor x.[15] Within each rating grade, obligors are statistically identical (except for loan size). That is, every obligor of grade ζ has unconditional default probability $\bar{p}_\zeta$ and has the same weight w_ζ on the systematic risk factor. (The value of w_ζ will, of course, depend on the choice of model.) The $\bar{p}$ values are set to the long-term average annual default probabilities given in Table 6.9 of the CreditMetrics Technical Document, and are shown below in the first column of Table 2. For a portfolio of loans, this is likely to be a relatively conservative calibration of mean annual default probilities.[16]

The weights w_ζ are calibrated for each model by working backwards from the historical volatility of annual default rates in each rating grade. First, using data in Brand and Bahar (1998, Table 12) on historical default experience in each grade, we estimate the variance V_ζ of the conditional default rate $p_\zeta(x)$. The estimation method and results are described in Appendix B. For calibration purposes, the default rate volatilities are most conveniently expressed as normalized standard deviations $\sqrt{V_\zeta}/\bar{p}_\zeta$. The values assumed in the simulations are shown in the second column of Table 2, and the implied default correlations ρ_ζ for any two obligors in the same rating grade are shown in the third column. To confirm the qualitative robustness of the results, additional simulations will be presented Section 4 in which the assumed normalized volatilities are twice the values used here.

The second step in determining the w_ζ is model-dependent. To calibrate the CreditMetrics weights, we use Proposition 1:

Proposition 1

In the CreditMetrics model,

$$V_\zeta \equiv \mathrm{Var}[p_\zeta(x)] = \mathrm{BIVNOR}(C_\zeta, C_\zeta, w_\zeta^2) - \bar{p}_\zeta^2 \tag{13}$$

---

[15]This is in the same spirit as the "Z-risk" approach of Belkin, Forest, Jr., Lawrence R. and Suchower (1998).

[16]Carey (1998) observes that default rates on speculative grade private placement loans tend to be lower than on publicly held bonds of the same senior unsecured rating, and attributes this superior performance to closer monitoring.

where BIVNOR(z_1, z_2, ρ) is the bivariate normal cdf for $Z \equiv [Z_1\ Z_2]'$ such that

$$\mathrm{E}[Z] = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \qquad \mathrm{Var}[Z] = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}.$$

The proof is given in Appendix C. Given the cut-off values C_ζ, which are functions of the $\bar{p}_\zeta$, and normalized volatilities $\sqrt{V_\zeta}/\bar{p}_\zeta$, nonnegative w_ζ are uniquely determined by the nonlinear equation (13). The solutions are shown in the fourth column of Table 2.

In CreditRisk$^+$, a model with a single systematic risk factor and obligor-specific idiosyncratic risk can be parameterized flexibly as a two risk factor model in which the first risk factor has zero volatility and thus always equals one.[17] Let w_ζ be the weights on x_2 (which are constant across obligors within a grade but allowed to vary across grades), so the weights on x_1 are $1 - w_\zeta$. To simplify notation, we set the first risk factor (x_1) identically equal to one, and denote the second risk factor (x_2) as x and the standard deviation σ_2 of x_2 as σ. Under this specification, the variance of the default probability in CreditRisk$^+$ for a grade ζ obligor is

$$V_\zeta = \mathrm{Var}[\bar{p}_\zeta(1 - w_\zeta + w_\zeta x)] = (\bar{p}_\zeta w_\zeta \sigma)^2$$

so the normalized volatility $\sqrt{V_\zeta}/\bar{p}_\zeta$ equals $w_\zeta \sigma$.

Given σ, the weights w_ζ are uniquely determined. However, there is no obvious additional information to bring to the choice of σ. This might appear to make little or no difference, because the volatility of the default probabilities depends only on the product $(w_\zeta \sigma)$. However, because σ controls the shape (and not merely the scale) of the distribution of x, higher moments of the distribution of $p_i(x)$ depend directly on σ and not only on the product $(w_\zeta \sigma)$. Consequently, tail probabilities for portfolio loss are quite sensitive to the choice of σ. To illustrate this sensitivity, simulation results will be presented for three values of σ (1.0, 1.5, and 4.0). See the last three columns of Table 2 for the values of w_ζ corresponding to each of these σ and the grade-specific normalized volatilities.

In the CreditRisk$^+$ manual, §A7.3, it is suggested that σ is roughly one. This estimate is based on a single-sector calibration of the model, which is equivalent to setting all the w_ζ to one. The exposure-weighted average of values in the second column of Table 2 would then be a reasonable calibration of σ. For most portfolios, this would yield $\sigma \approx 1$, as suggested. Our specification is strictly more general than the single-sector approach, because it allows the relative importance of idiosyncratic risk to vary across rating grades, and also is more directly comparable to our calibrated correlation structure for CreditMetrics. Note that the difficulty of calibrating σ in this more general specification should not be interpreted as a disadvantage to CreditRisk$^+$ relative to CreditMetrics, because CreditMetrics avoids this calibration issue by fiat. In assuming the normal distribution for the systematic risk factor, CreditMetrics is indeed imposing very strong restrictions

---

[17]See the CreditRisk$^+$ manual, §A12.3. The first factor is referred to as a "specific factor." Because it represents diversifiable risk, it contributes no volatility to a well-diversified portfolio.

Table 2: Default Rate Volatility and Factor Weights[*]

	Historical Experience			Systematic Risk Weights			
	$\bar{p}$	$\sqrt{V}/\bar{p}$	ρ	CM2S	CR+	CR+	CR+
σ					1.0	1.5	4.0
AAA	0.01	1.4	0.0002	0.272	1.400	0.933	0.350
AA	0.02	1.4	0.0004	0.285	1.400	0.933	0.350
A	0.06	1.2	0.0009	0.279	1.200	0.800	0.300
BBB	0.18	0.4	0.0003	0.121	0.400	0.267	0.100
BB	1.06	1.1	0.0130	0.354	1.100	0.733	0.275
B	4.94	0.55	0.0157	0.255	0.550	0.367	0.138
CCC	19.14	0.4	0.0379	0.277	0.400	0.267	0.100

*: Unconditional annual default probabilities $\bar{p}$ taken from the Credit-
    Metrics Technical Document, Table 6.9., and are expressed here in
    percentage points. Historical experience for default rate volatility
    derived from Brand and Bahar (1998, Table 12), as described below
    in Appendix B.

on the shape of the distribution tail.

When $\sigma = 1$ is used in our calibration of CreditRisk+, a problem arises in that some of the systematic risk weights exceed one. Such values imply negative weights on the specific factors, which violate both intuition and the formal assumptions of the model. However, CreditRisk+ can in principle tolerate negative weights so long as all coefficients in the polynomial expansion of the portfolio loss probability generating function remain positive.[18] For the weights in the $\sigma = 1$ column, we have confirmed numerically that our simulations always produce valid loss distributions.

3.3 Main simulation results

Results for the main set of simulations are displayed in Table 3.[19] Each quadrant of the table shows summary statistics and selected percentile values for CreditMetrics and CreditRisk+ portfolio loss distributions for a portfolio of a given credit quality distribution. The summary statistics are the mean, standard deviation, index of skewness and index of kurtosis. The latter two are defined for

---

[18]Conditional on small realizations of x, an obligor with negative weight on the specific factor can have a negative default probability. However, so long as such obligors are relatively few and their negative weights relatively small in magnitude, the portfolio loss distribution can still be well-behaved. There may be some similarity to the problem of generating default probabilities *over one* conditional on large realizations of x, which need not cause any problem at the portfolio level, so long as the portfolio does not have too many low-rated obligors with high loading on the systematic risk factor.

[19]In these simulations, there are $N = 5000$ loans in the portfolio, grade-specific loan size distributions are taken from the SoA sample, average severity of loss is held constant at 30%, the weights w_ζ and CreditRisk+ parameter σ are taken from Table 2. CreditMetrics distributions are formed using 200,000 portfolio draws.

a random variable y by

$$\text{Skewness}(y) = \frac{\text{E}[(y - \text{E}[y])^3]}{\text{Var}[y]^{3/2}}, \qquad \text{Kurtosis}(y) = \frac{\text{E}[(y - \text{E}[y])^4]}{\text{Var}[y]^2}.$$

Skewness is a measure of the asymmetry of a distribution, and kurtosis is a measure of the relative thickness of the tails of the distribution. For portfolio credit risk models, high kurtosis indicates a relatively high probability of very large credit losses.

The percentile values presented in the table are the loss levels associated with the 50% (median), 75%, 95%, 99%, 99.5% and 99.97% points on the cumulative distribution of portfolio losses. In many discussions of credit risk modeling, the 99th and sometimes the 95th percentiles of the distribution are taken as points of special interest. The 99.5th and 99.97th percentiles may appear to be extreme tail values, but are in fact of greater practical interest than the 99th percentile. To merit a AA rating, an institution must have a probability of default over a one year horizon of roughly three basis points (0.03%).[20] Such an institution therefore ought to hold capital (or reserves) against credit loss equal to the 99.97th percentile value. Capitalization sufficient to absorb up to the 99.5th percentile value of losses would be consistent with only a BBB− rating.

Table 3, for the Average quality portfolio, illustrates the qualitative characteristics of the main results. The expected loss under either model is roughly 48 basis points of the portfolio book value.[21] The standard deviation of loss is roughly 32 basis points. When the CreditRisk[+] parameter σ is set to 1, the two models predict roughly similar loss distributions overall. The 99.5th and 99.97th percentile values are roughly 1.8% and 2.7% of portfolio book value in each case. As σ increases, however, the CreditRisk[+] distribution becomes increasingly kurtotic. The standard deviation remains roughly the same, but tail percentile values increase substantially. The 99.5th and 99.97th CreditRisk[+] percentile values given $\sigma = 4.0$ are respectively 40% and 90% larger than the corresponding CreditMetrics values.

High, Low, and Very Low quality portfolios produce different expected losses (19, 93, and 111 basis points, respectively), but similar overall conclusions regarding our comparison of the two models. CreditRisk[+] with $\sigma = 1.0$ produces distributions roughly similar to those of CreditMetrics, although as credit quality deteriorates the extreme percentile values in CreditRisk[+] increase more quickly than in CreditMetrics. As σ increases, so do the extreme loss percentiles.

Overall, capital requirements implied by these simulations may seem relatively low. Even with a Low quality portfolio, a bank would need to hold only 4.5%-6% capital against credit risk in order to maintain a AA rating standard.[22] It should be noted, however, that these simulations assume uniform default correlations within each rating grade. In real world portfolios, there may

---

[20]This is a rule of thumb often used by practioners. Following the CreditMetrics Technical Document, we have taken a slightly lower value (0.02%) as the AA default probability.

[21]For this credit quality distribution, the expected annual default rate is 1.6% (by loan value). Multiply by the average severity of 30% to get a loss of 48 basis points.

[22]Simulations by Carey (1998) suggest somewhat higher capital requirements. His simulations account for recovery risk, which is assumed away here. Perhaps more importantly, his simulations are calibrated using data from 1986–92, which was a relatively unfavorable period in the credit cycle.

sometimes be pockets of higher default correlation, due perhaps to imperfect geographic or industry diversification. Furthermore, it should be emphasized that these simulations incorporate only default risk, and thus additional capital must be held for other forms of risk, including market risk, operational risk, and recovery uncertainty.

Table 3: CreditMetrics vs CreditRisk$^+$: Main Simulations

	High Quality Portfolio				Average Quality Portfolio			
	CM2S	CR+			CM2S	CR+		
σ		1.00	1.50	4.00		1.00	1.50	4.00
Mean	0.194	0.194	0.194	0.194	0.481	0.480	0.480	0.480
Std Dev	0.152	0.155	0.155	0.155	0.319	0.325	0.324	0.323
Skewness	1.959	1.864	2.515	5.694	1.696	1.844	2.611	6.360
Kurtosis	9.743	8.285	13.157	59.539	8.137	8.228	13.823	69.379
0.5000	0.156	0.150	0.148	0.160	0.409	0.391	0.384	0.414
0.7500	0.257	0.257	0.240	0.222	0.624	0.612	0.567	0.520
0.9500	0.486	0.501	0.497	0.398	1.089	1.120	1.116	0.869
0.9900	0.733	0.745	0.794	0.858	1.578	1.628	1.749	1.916
0.9950	0.847	0.850	0.928	1.121	1.795	1.847	2.033	2.488
0.9997	1.342	1.277	1.490	2.345	2.714	2.736	3.225	5.149

	Low Quality Portfolio				Very Low Quality Portfolio			
	CM2S	CR+			CM2S	CR+		
σ		1.00	1.50	4.00		1.00	1.50	4.00
Mean	0.927	0.927	0.927	0.917	1.107	1.106	1.106	1.106
Std Dev	0.557	0.565	0.565	0.486	0.635	0.644	0.643	0.641
Skewness	1.486	1.872	2.711	4.898	1.393	1.874	2.724	6.883
Kurtosis	6.771	8.362	14.455	35.767	6.299	8.374	14.540	77.125
0.5000	0.809	0.769	0.753	0.815	0.977	0.926	0.906	0.979
0.7500	1.194	1.154	1.063	0.967	1.418	1.364	1.259	1.146
0.9500	1.989	2.045	2.041	1.585	2.316	2.379	2.376	1.854
0.9900	2.782	2.936	3.161	3.481	3.187	3.395	3.654	4.024
0.9950	3.124	3.320	3.664	4.504	3.562	3.832	4.227	5.192
0.9997	4.558	4.877	5.770	9.251	5.105	5.607	6.631	10.618

4   Robustness of model results

In this section, we explore the sensitivity of the models to parameter calibration and portfolio construction.

Obligor count: Compared to portfolios of equities, loan portfolios can be quite large and still receive substantial diversification benefits from adding more obligors.[23] Table 4 compares Credit-Metrics and CreditRisk$^+$ results for Average quality portfolios of 1000, 5000, and 10,000 obligors. Even with portfolios of this size, increasing the number of obligors reduces risk significantly. The standard deviation of the 10,000 obligor portfolio is roughly 20% less than that of the 1000 obligor portfolio, and the 99.5th and 99.97th percentile values fall by 13–15%. However, the qualitative nature of the results, particularly the comparison between the two models, remains unchanged.

Table 4: Effect of Obligor Count on Portfolio Loss Distributions*

	$N = 1000$		$N = 5000$		$N = 10,000$	
	CM2S	CR+	CM2S	CR+	CM2S	CR+
Mean	0.480	0.480	0.481	0.480	0.480	0.480
Std Dev	0.387	0.398	0.319	0.324	0.306	0.312
Skewness	1.672	2.226	1.696	2.611	1.734	2.764
Kurtosis	7.442	11.133	8.137	13.823	8.390	14.774
0.5000	0.383	0.370	0.409	0.384	0.410	0.380
0.7500	0.653	0.619	0.624	0.567	0.615	0.552
0.9500	1.235	1.251	1.089	1.116	1.064	1.097
0.9900	1.803	1.957	1.578	1.749	1.531	1.719
0.9950	2.044	2.278	1.795	2.033	1.750	1.999
0.9997	3.093	3.626	2.714	3.225	2.653	3.169

*: Average quality portfolio with SoA loan size distributions. All CreditRisk$^+$ simulations use $\sigma = 1.5$.

Loan size distribution: The loan size distributions derived from the SoA data are likely to be somewhat skew in comparison with real bank portfolios. For $N = 5000$, the largest loans are over 0.65% of the portfolio, which is not much below supervisory concentration limits. To examine the effect of loan size distribution, we construct portfolios in which all loans within a single rating grade have the same size. Results are shown in Table 5 for Average quality portfolios. The tail percentiles are somewhat lower for the equal-sized portfolios, but, if one considers the magnitude of differences between the two loan-size distributions, the difference in model outputs seems minor.

---

[23]Essentially, this is because risk in loans is dominated by large changes in value which occur with relatively low probability events. The skew distribution of individual losses allow the tail of the portfolio loss distribution to thin with diversification at only a relatively slow rate.

These results suggest that, with real bank portfolios, neither model is especially sensitive to the distribution of loan sizes.

Table 5: Equal-sized vs SoA Loan Sizes*

| | SoA Loan Sizes | | Equal-Sized Loans | |
	CM2S	CR+	CM2S	CR+
Mean	0.481	0.480	0.481	0.486
Std Dev	0.319	0.324	0.299	0.308
Skewness	1.696	2.611	1.801	2.901
Kurtosis	8.137	13.823	8.712	15.597
0.5000	0.409	0.384	0.412	0.373
0.7500	0.624	0.567	0.609	0.549
0.9500	1.089	1.116	1.051	1.101
0.9900	1.578	1.749	1.527	1.728
0.9950	1.795	2.033	1.747	2.009
0.9997	2.714	3.225	2.649	3.187

*: Average quality portfolio of $N = 5000$ obligors.
CreditRisk$^+$ simulations use $\sigma = 1.5$.

Normalized volatilities: Due to the empirical difficulty of estimating default correlations with precision, practioners may be especially concerned with the sensitivity of the results to the values of the normalized volatilities in Table 2. Therefore, we calibrate and run a set of simulations in which normalized volatilities are double the values used above. CreditMetrics weights w_ζ increase substantially, though not quite proportionately.[24] We retain the same CreditRisk$^+$ w_ζ values given in the last three columns of Table 2, but double the respective σ values.

Results are presented in Table 6 for the Average quality portfolio. As should be expected, extreme tail percentile values increase substantially. Compared to the values in Table 3, the 99.97th percentile values nearly triple. Similar increases in tail percentile values are observed for the other credit quality distributions.

Discretization of loan sizes: In the main simulations, the CreditRisk$^+$ base exposure unit is set to λ times the fifth percentile value of the distribution of loan sizes. At least locally, the error introduced by this discretization is negligible. We have run most of our simulations with ν_0 set to λ times the 2.5th and 10th percentile values. For both these alternatives, the percentile values of the loss distribution differed from those of the main simulations by no more than 0.0005.

---

[24]Due to the nonlinearity of the normal cdf, a given percentage increase in the normalized volatility is generally associated with a somewhat smaller percentage increase in the weight on x.

19

Table 6: Effect of Increased Default Volatilities*

σ	CM2S	CR+		
		2.00	3.00	8.00
Mean	0.480	0.480	0.480	0.480
Std Dev	0.590	0.610	0.609	0.603
Skewness	3.221	3.860	5.673	14.334
Kurtosis	20.278	25.182	50.960	301.953
0.5000	0.287	0.265	0.313	0.400
0.7500	0.615	0.507	0.447	0.492
0.9500	1.597	1.648	1.442	0.710
0.9900	2.845	3.130	3.311	2.533
0.9950	3.467	3.818	4.239	4.202
0.9997	6.204	6.772	8.386	13.459

*: Average quality portfolio with SoA loan size distributions.
$N = 5000$ obligors.

5 Modified CreditRisk$^+$ specifications

The analysis of Section 3.3 demonstrates the sensitivity of CreditRisk$^+$ to the calibration of σ. When we vary σ while holding the $w_\zeta \sigma$ constant, the mean and standard deviation of loss remain approximately unchanged, but the tail percentile values change markedly.[25] This sensitivity is both a direct and an indirect consequence of a property of the gamma distribution for x. Unlike the normal distribution, which has kurtosis equal to 3 regardless of its variance, the kurtosis of a gamma-distributed variable depends on its parameters. A gamma random variable with mean one and variance σ^2 has kurtosis $3(1 + 2\sigma^2)$, so higher σ imposes a more fat-tailed shape on the distribution, which is transmitted to the shape of the distribution for $p_i(x)$ for each obligor i. So long as $w_{\zeta(i)}\sigma$ is held constant, varying σ has no effect on the mean or standard deviation of $p_i(x)$. However, it is straightforward to show that the kurtosis of $p_i(x)$ equals the kurtosis of x, so increasing σ increases the kurtosis of $p_i(x)$.

Increasing the kurtosis of the $p_i(x)$ has the direct effect of increasing the thickness of the tail of the distribution for loss. This is explored below by substituting an alternative distribution for x which has mean one and variance σ^2 but is less kurtotic. The indirect effect of higher kurtosis for $p_i(x)$ is that it magnifies the error induced by the Poisson approximation. To explore the effect of σ on the size of the approximation error, we use the methods of Section 2.2 to eliminate the Poisson approximation from the calculations.

The Poisson approximation necessarily contributes to the thickness of the tail in CreditRisk$^+$ loss distributions. In a Monte Carlo based model, such as CreditMetrics, an obligor can default no more than once, so no more than N defaults can be suffered. Under the Poisson approximation, a

---

[25]They change slightly because higher σ increases the frequency with which $p_i(x)$ exceeds one, which distorts the analytical results.

single obligor can be counted in default any number of times (albeit with very small probabilities of multiple defaults). Thus, CreditRisk$^+$ assigns a positive probability on the number of defaults exceeding the number of obligors. No matter how the portfolio is constructed and how the two models are calibrated, there must be a crossing point beyond which CreditRisk$^+$ percentile values all exceed the corresponding CreditMetrics percentile values.

Depending on the portfolio and the model parameters, the effect of the Poisson approximation may or may not be negligible. To test the empirical relevance of this effect, we compare CreditRisk$^+$ results against those of a Monte Carlo version of CreditRisk$^+$. The Monte Carlo version is similar to that outlined in Section 2.2, except that default of obligor i, conditional on x, is drawn as a Bernoulli random variate with probability $p_i(x)$, rather than using the latent variable approach of equation (10). This avoids the small approximation error induced by equation (11), but otherwise imposes exactly the same distributional assumptions and functional forms as the standard CreditRisk$^+$ model.

We conduct similar Monte Carlo exercises to explore alternative distributional assumptions for x. Say that x is distributed such that $x^2 \sim \text{Gamma}(\alpha, \beta)$. As described in Appendix D, it is straightforward to solve for parameters (α, β) such that the $\text{E}[x] = 1$ and $\text{Var}[x] = \sigma^2$. Although this x matches the mean and variance of the standard CreditRisk$^+$ gamma-distributed risk factor, it is much less kurtotic. The "gamma-squared" distribution is compared to the ordinary gamma distribution in Figure 2. The top panel plots the cdfs for a gamma distributed variable (solid line) and a gamma-squared distributed variable (dashed line). Both variables have mean one and variance one. The two distributions appear to be quite similar, and indeed would be difficult to distinguish empirically. Nonetheless, as shown in the bottom panel, the two distributions differ substantially in the tails. The 99.9th percentiles are 6.91 and 5.58 for the gamma and gamma-squared distributions, respectively. The 99.97th percentiles are 8.11 vs. 6.20.

The results of both exercises on an Average quality portfolio are shown in Table 7.[26] The standard CreditRisk$^+$ results for $\sigma = 1.5$ and $\sigma = 4.0$ (columns 1 and 4) are taken from Table 3. Results for the Monte Carlo version of CreditRisk$^+$ are shown in columns 2 and 5. For the moderate value of $\sigma = 1.5$, the 99.97th percentile value is reduced by under two percent. For the larger value $\sigma = 4.0$, however, the 99.97th percentile value is reduced by over eight percent. The higher the value of σ, the higher the probability of large conditional default probabilities. As the validity of the Poisson approximation thus breaks down for high σ, so does the accuracy of the analytic CreditRisk$^+$ methodology.

Results for the modified CreditRisk$^+$ with x^2 gamma-distributed are shown in columns 3 and 6. For both values of σ, the mean and standard deviation of portfolio loss are roughly as before, but the tail percentiles are quite significantly reduced. Indeed, the 99.97th percentile value for the modified model under $\sigma = 1.5$ is even less than the corresponding CreditMetrics value. This demonstrates the critical importance of the shape of the distribution of the systematic risk factor.

---

[26]Qualitatively similar results are found for portfolios based on the other credit quality distributions.

Figure 2: Gamma and Gamma-Squared Distributions
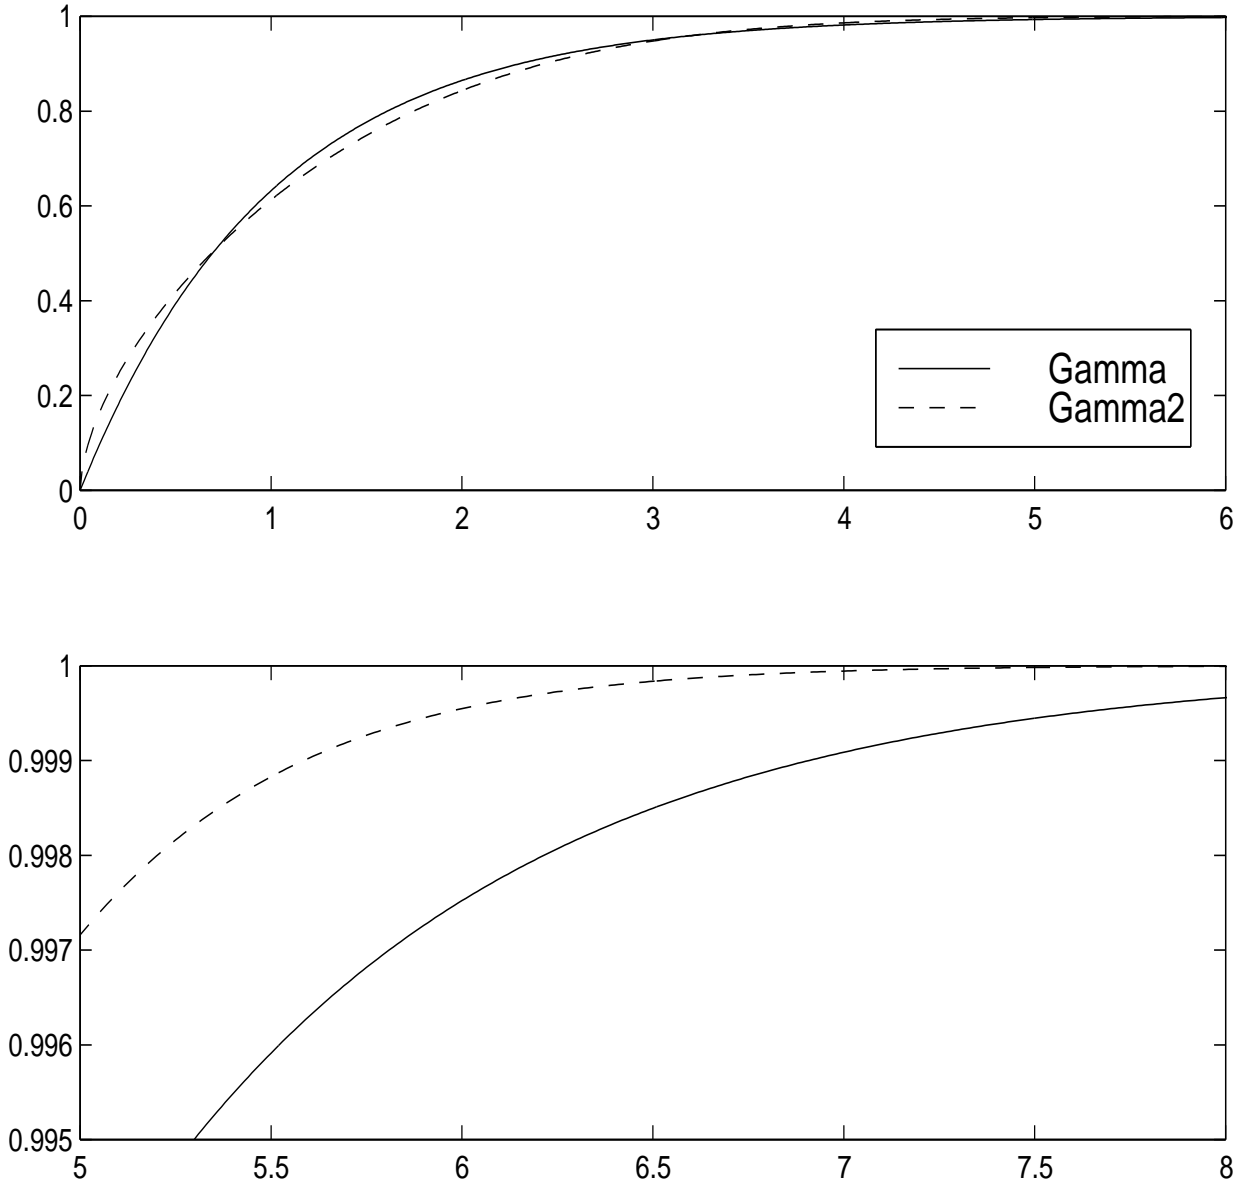
Note: The two lines are cdfs of variables with mean one and variance one. If x is gamma distributed, the solid line is its cdf. If x^2 is gamma distributed, the dashed line is its cdf.

Table 7: Modified CreditRisk$^+$ Models^{a,b}

| | $\sigma = 1.50$ | | | $\sigma = 4.00$ | | |
	CR+	CR+(MC)	CR+(X2)	CR+	CR+(MC)	CR+(X2)
Mean	0.480	0.479	0.480	0.480	0.480	0.479
Std Dev	0.324	0.319	0.322	0.323	0.316	0.321
Skewness	2.611	2.606	2.037	6.360	6.070	5.142
Kurtosis	13.823	13.938	8.342	69.379	62.525	40.076
0.5000	0.384	0.385	0.374	0.414	0.416	0.411
0.7500	0.567	0.564	0.575	0.520	0.518	0.512
0.9500	1.116	1.106	1.158	0.869	0.866	0.888
0.9900	1.749	1.719	1.670	1.916	1.904	2.047
0.9950	2.033	1.991	1.868	2.488	2.494	2.553
0.9997	3.225	3.179	2.561	5.149	4.729	4.099

a: Average quality portfolio with SoA loan size distributions. $N = 5000$ obligors. 200,000 portfolio draws in the Monte Carlo simulations.

b: CR+ (columns 1 and 4) is standard CreditRisk$^+$. CR+(MC) (columns 2 and 5) is CreditRisk$^+$ estimated by Monte Carlo. CR+(X2) (columns 3 and 6) is a Monte Carlo CreditRisk$^+$ with x^2 gamma-distributed.

Discussion

This paper demonstrates that there is no unbridgeable difference in the views of portfolio credit risk embodied in the two models. If we consider the restricted form of CreditMetrics used in the analysis, then each model can be mapped into the mathematical framework of the other, so that the primary sources of discrepancy in results are differences in distributional assumptions and functional forms.

Simulations are constructed for a wide range of plausible loan portfolios and correlation parameters. The results suggest a number of general conclusions. First, the two models perform very similarly on an average quality commercial loan portfolio when the CreditRisk$^+$ volatility parameter σ is given a low value. Both models demand higher capital on lower quality portfolios, but CreditRisk$^+$ is somewhat more sensitive to credit quality than the two-state version of CreditMetrics. It should be emphasized, however, that the full implementation of CreditMetrics encompasses a broader notion of credit risk, and is likely to produce somewhat larger tail percentiles than our restricted version.

Second, results do not depend very strongly on the distribution of loan sizes within the portfolio, at least within the range of size concentration normally observed in bank portfolios. The discretization of loan sizes in CreditRisk$^+$ has negligible impact.

Third, both models are highly sensitive to the volatility of default probabilities, or, equivalently, to the average default correlations in the portfolio. When the standard deviation of the default probabilities is doubled, required capital increases by two to three times.

Finally, the models are highly sensitive to the shape of the implied distribution for default probabilities. CreditMetrics, which implies a relatively thin-tailed distribution, reports relatively low tail percentile values for portfolio loss. The tail of CreditRisk$^+$ depends strongly on the parameter σ, which determines the kurtosis (but not the mean or variance) of the distribution of default probabilities. Choosing less kurtotic alternatives for the gamma distribution used in CreditRisk$^+$ sharply reduces its tail percentile values for loss without affecting the mean and variance.

This sensitivity ought to be of primary concern to practioners. It is difficult enough to measure expected default probabilities and their volatility. Capital decisions, however, depend on extreme tail percentile values of the loss distribution, which in turn depend on higher moments of the distribution of default probabilities. These higher moments cannot be estimated with any precision given available data. Thus, the models are more likely to provide reliable measures for comparing the relative levels of risk in two portfolios than to establish authoritatively absolute levels of capital required for any given portfolio.

A   Properties of the gamma distribution

The gamma distribution is a two parameter distribution commonly used in time-to-failure and other engineering applications. If x is distributed Gamma(α, β), the probability density function of x is given by

$$f(x|\alpha, \beta) = \frac{x^{\alpha-1}\exp(-x/\beta)}{\beta^\alpha \Gamma(\alpha)} \tag{14}$$

where $\Gamma(\alpha)$ is the Gamma function.[27] The mean and variance of x are given by $\alpha\beta$ and $\alpha\beta^2$, respectively. Therefore, if we impose $\mathrm{E}[x] = 1$ and $\mathrm{V}[x] = \sigma^2$, then we must have $\alpha = 1/\sigma^2$ and $\beta = \sigma^2$.

B   Estimating the volatility of default probabilities

This appendix demonstrates a simple nonparametric method of estimating the volatility of default probabilities from historical performance data published by Standard & Poor's (Brand and Bahar 1998, Table 12). Let $p_\zeta(x_t)$ denote the probability of default of a grade ζ obligor, conditional on the realized value x_t of a systematic risk factor. We need to estimate the unconditional variance $\mathrm{V}[p_\zeta(x)]$. We assume that the x_t are serially independent and that obligor defaults are independent conditional on x_t. Both CreditMetrics and CreditRisk$^+$ satisfy this framework, though the two models impose different distributional assumptions for x and functional forms for $p(x)$.

For each year in 1981–97 and for each rating grade, S&P reports the number of corporate obligors in its ratings universe on January 1, and the number of obligors who have defaulted by

---

[27]Other parameterizations of this distribution are sometimes seen in the literature. This is the parameterization used by Credit Suisse Financial Products (1997, Eq. 50).

the end of the calendar year. Let $\hat{d}_{\zeta t}$ be the number of grade ζ defaults during year t, and let $\hat{n}_{\zeta t}$ denote the number of grade ζ obligors at the start of year t. Let $\hat{p}_{\zeta t}$ denote the observed default frequency $\hat{d}_{\zeta t}/\hat{n}_{\zeta t}$. We assume that the size of the universe $\hat{n}_{\zeta t}$ is independent of the realization of x_t.

The general rule for conditional variance is

$$\mathrm{V}[y] = \mathrm{E}[\mathrm{V}[y|z]] + \mathrm{V}[\mathrm{E}[y|z]]. \tag{15}$$

Applied to the problem at hand, we have

$$\mathrm{V}[\hat{p}_\zeta] = \mathrm{E}[\mathrm{V}[\hat{p}_\zeta|p_\zeta(x), \hat{n}_\zeta]] + \mathrm{V}[\mathrm{E}[\hat{p}_\zeta|p_\zeta(x), \hat{n}_\zeta]] \tag{16}$$

Obligor defaults are independent conditional on x, so $\hat{d}_{\zeta t} \sim \mathrm{Binomial}(\hat{n}_{\zeta t}, p_\zeta(x_t))$. The expectation of the conditional variance of $\hat{p}_{\zeta t}$ is therefore given by

$$
\begin{aligned}
\mathrm{E}[\mathrm{V}[\hat{p}_\zeta|p_\zeta(x), \hat{n}_\zeta]] &= \mathrm{E}[\mathrm{V}[\hat{d}_\zeta|p_\zeta(x)]/\hat{n}_\zeta^2] = \mathrm{E}[p_\zeta(x)(1 - p_\zeta(x))/\hat{n}_\zeta] \\
&= \mathrm{E}[1/\hat{n}_\zeta](\mathrm{E}[p_\zeta(x)] - (\mathrm{V}[p_\zeta(x)] + \mathrm{E}[p_\zeta(x)]^2)) = \mathrm{E}[1/\hat{n}_\zeta](\bar{p}_\zeta(1 - \bar{p}_\zeta) - \mathrm{V}[p_\zeta(x)])
\end{aligned}
$$

where the second equality follows from the formula for the variance of a binomial random variable; the third equality follows from the mutual independence of x and $\hat{n}_\zeta$ and the rule $\mathrm{V}[y] = \mathrm{E}[y^2] - \mathrm{E}[y]^2$; and the final equality from $\mathrm{E}[p_\zeta(x)] = \bar{p}_\zeta$.

Since $\mathrm{E}[\hat{p}_\zeta|p_\zeta(x), \hat{n}_\zeta] = p_\zeta(x)$, the last term in equation (16) is simply $\mathrm{V}[p_\zeta(x)]$. Substitute these simplified expressions into equation (16) and rearrange to obtain

$$\mathrm{V}[p_\zeta(x)] = \frac{\mathrm{V}[\hat{p}_\zeta] - \mathrm{E}[1/\hat{n}_\zeta]\bar{p}_\zeta(1 - \bar{p}_\zeta)}{1 - \mathrm{E}[1/\hat{n}_\zeta]}. \tag{17}$$

The values of $\bar{p}_\zeta$ observed in the S&P data differ slightly from the values used for calibration in Section 3.2. It is most convenient, therefore, to normalize the estimated default rate volatilities as ratios of the standard deviation of $p_\zeta(x)$ to its expected value, $\sqrt{V[p_\zeta(x)]}/\bar{p}_\zeta$. In the first two columns of Table 8, we present for each rating grade the empirical values of $\bar{p}$ and $\mathrm{E}[1/\hat{n}]$ in the S&P data. The third column presents the observed variance of default rates, $\hat{V}[\hat{p}]$, expressed in normalized form. The fourth column gives the implied normalized volatilities for the unobserved true conditional default probabilities.

For the highest grades, AAA and AA, no defaults occurred in the S&P sample, so it is impossible to estimate a volatility for these grades. Among the A obligors, only five defaults were observed in the sample, so the default volatility is undoubtedly measured with considerable imprecision. Therefore, calibration of the normalized volatilities for these grades requires some judgement. Our chosen values for these ratios are given in the final column of Table 8. It is assumed that normalized volatilities are somewhat higher for the top grades, but that the estimated value for grade A is implausibly high. For the lower grades, the empirical estimates are made with greater precision

(due to the larger number of defaults in sample), so these values are maintained.

Table 8: Empirical Default Frequency and Volatility

	$\bar{p}_\zeta$	$\mathrm{E}[1/\hat{n}_\zeta]$	$\sqrt{\hat{V}[\hat{p}_\zeta]}/\bar{p}_\zeta$	$\sqrt{\hat{V}[p_\zeta(x)]}/\bar{p}_\zeta$	$\sqrt{V_\zeta(x)}/\bar{p}_\zeta$
AAA	0	0.0092	.	.	1.4
AA	0	0.0030	.	.	1.4
A	0.0005	0.0017	2.4857	1.5896	1.2
BBB	0.0018	0.0026	1.2477	0.3427	0.4
BB	0.0091	0.0038	1.2820	1.1108	1.1
B	0.0474	0.0041	0.6184	0.5492	0.55
CCC	0.1890	0.0360	0.5519	0.3945	0.4

C Proof of Proposition 1

Let y_1 and y_2 be the CreditMetrics latent variables for two grade ζ obligors. Assume that there is only one systematic risk factor and that the two obligors have the same weight w_ζ on that risk factor. Thus,

$$
\begin{aligned}
y_1 &= xw_\zeta + \sqrt{1 - w_\zeta^2}\,\epsilon_1 \\
y_2 &= xw_\zeta + \sqrt{1 - w_\zeta^2}\,\epsilon_2.
\end{aligned}
$$

Conditional on x, default events for these obligors are independent, so

$$\Pr(y_1 < C_\zeta \,\&\, y_2 < C_\zeta | x) = \Pr(y_1 < C_\zeta | x)\Pr(y_2 < C_\zeta | x) = \Phi\left((C_\zeta - xw_\zeta)/\sqrt{1 - w_\zeta^2}\right)^2 = p_\zeta(x)^2.$$

Therefore,

$$\mathrm{Var}[p_\zeta(x)] = \mathrm{E}[p_\zeta(x)^2] - \mathrm{E}[p_\zeta(x)]^2 = \mathrm{E}[\Pr(y_1 < C_\zeta \,\&\, y_2 < C_\zeta | x)] - \mathrm{E}[p_\zeta(x)]^2.$$

Since y_1 and y_2 each have mean zero and variance one, and have correlation w_ζ^2, the unconditional expectation $\mathrm{E}[\Pr(y_1 < C_\zeta \,\&\, y_2 < C_\zeta | x)]$ is given by $\mathrm{BIVNOR}(C_\zeta, C_\zeta, w_\zeta^2)$. This gives

$$\mathrm{Var}[p_\zeta(x)] = \mathrm{BIVNOR}(C_\zeta, C_\zeta, w_\zeta^2) - \bar{p}_\zeta^2,$$

as required.

D   Distribution of an exponentiated gamma random variable

Assume that x is distributed such that $x^{1/r} \sim \text{Gamma}(\alpha, \beta)$ for some given $r > 0$. We wish to solve for parameters (α, β) such that $\text{E}[x] = 1$ and $\text{Var}[x] = \sigma^2$.

It is straightforward to show that the moments of x are given by

$$\text{E}[x^k] = \frac{\Gamma(\alpha + rk)}{\Gamma(\alpha)} \beta^{rk}.$$

Therefore,

$$\text{E}[x] = \frac{\Gamma(\alpha + r)}{\Gamma(\alpha)} \beta^r = 1 \quad \Rightarrow \quad \beta = \left(\frac{\Gamma(\alpha)}{\Gamma(\alpha + r)}\right)^{1/r}.$$

To solve for α, use the variance restriction

$$\text{Var}[x] = \text{E}[x^2] - \text{E}[x]^2 = \sigma^2$$

which implies

$$
\begin{aligned}
\sigma^2 + 1 = \text{E}[x^2] &= \frac{\Gamma(\alpha + 2r)}{\Gamma(\alpha)} \beta^{2r} \\
&= \frac{\Gamma(\alpha + 2r)}{\Gamma(\alpha)} \left(\frac{\Gamma(\alpha)}{\Gamma(\alpha + r)}\right)^2 = \frac{\Gamma(\alpha + 2r)\Gamma(\alpha)}{\Gamma(\alpha + r)^2}.
\end{aligned}
\tag{18}
$$

For any $r > 0$, there exists a unique solution to this equation for α. To see this, define $g_r(\alpha)$ as the right hand side of equation (18). The gamma function is continuous for non-negative arguments, so $g_r(\alpha)$ is continuous as well. Note that $\Gamma(0) = \infty$, so $g_r(0) = \infty$ for $r > 0$. Using the infinite product form for the gamma function, given in Abramowitz and Stegun, eds (1968, 6.1.3), it is straightforward to show that

$$g_r(\alpha) = \prod_{j=0}^{\infty} \frac{(\alpha + r + j)^2}{(\alpha + j)(\alpha + 2r + j)}.$$

The limit of this expression as $\alpha \to \infty$ is one. Therefore, by the Intermediate Value Theorem, there exists a unique solution to $g_r(\alpha) = 1 + \sigma^2$.

References

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

Belkin, Barry, Forest, Jr., Lawrence R., and Stephan Suchower, "The Effect of Systematic Credit Risk on Loan Portfolio Value-at-Risk and Loan Pricing," *CreditMetrics Monitor*, 1998, *First Quarter*, 17–28.

Brand, Leo and Reza Bahar, "Ratings Performance 1997: Stability & Transition," Special Report, Standard & Poor's 1998.

Carey, Mark, "Credit Risk in Private Debt Portfolios," *Journal of Finance*, 1998, *10* (10), 56–61.

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

Finger, Christopher C., "Credit Derivatives in CreditMetrics," *CreditMetrics Monitor*, 1998, *Third Quarter*, 13–27.

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

Johnson, Norman L. and Samuel Kotz, *Distributions in Statistics: Discrete Distributions*, Houghton Mifflin, 1969.

Society of Actuaries, *1986–92 Credit Risk Loss Experience Study: Private Placement Bonds*, Schaumberg, IL: Society of Actuaries, 1996.