From CreditMetrics to CreditRisk$^+$ and Back Again

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

June 23, 1998

Abstract

In the short time since their public releases in 1997, J.P. Morgan's CreditMetrics and Credit Suisse's CreditRisk$^+$ have become influential benchmarks for internal credit risk models. Practitioners and policy makers have invested in implementing and exploring each of the models individually, but have made less progress with comparative analyses. Direct comparison of the models is not straightforward, because the two models are presented within rather different mathematical frameworks. One is familiar to econometricians as an ordered probit model, the other is based on insurance industry models of event risk. CreditMetrics and CreditRisk$^+$ may be addressing the same topic, but they appear to speak in different languages.

This paper develops methods for translating between these two languages. I show how a restricted version of CreditMetrics can be run through the mathematical machinery of CreditRisk$^+$, and how CreditRisk$^+$ can be mapped into a version of CreditMetrics. A series of simulation exercises uses these translation methods to evaluate the robustness of each model to the assumptions of the other, and to isolate the models' most sensitive restrictions.

*The views expressed herein are my own and do not necessarily reflect those of the Board of Governors or its staff. I am grateful for the helpful comments of Mark Carey and David Jones. Please address correspondence to the author at Division of Research and Statistics, Mail Stop 153, Federal Reserve Board, Washington, DC 20551, USA. Phone: (202)452-3705. Fax: (202)452-5295. Email: ⟨mgordy@frb.gov⟩.

In the short time since their public releases in 1997, J.P. Morgan's CreditMetrics and Credit Suisse's CreditRisk$^+$ have become influential benchmarks for internal credit risk models. 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. 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 of default 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 be addressing the same topic, but they appear to speak in different languages.

The purpose of this paper is to show how to translate between these two languages. Section 1 shows how a restricted version of CreditMetrics can be run through the mathematical machinery of CreditRisk$^+$. Section 2 maps CreditRisk$^+$ into a version of CreditMetrics. Section 3 (*not* *yet written*) uses these translation methods to develop comparative simulations. The goal of the simulation exercises is to isolate the models' most sensitive assumptions. Concluding remarks follow.

NOTE: The remainder of this draft assumes technical familiarity with both models. The final version will provide a somewhat thicker cushion of introductory material. Note also that my choice of notation below follows neither model exactly. My aim is to ease comparison across the two models, rather than to minimize distance to the models' original notation.

1 Mapping CreditMetrics to the CreditRisk$^+$ framework

CreditRisk$^+$ is essentially 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*. CreditMetrics is strictly more general in that it models migrations among multiple states (i.e., credit ratings) and allows for idiosyncratic uncertainty in loss given default. To map CreditMetrics into the CreditRisk$^+$ framework, form a restricted version of CreditMetrics with two states and fixed exposures (i.e., non-stochastic loss given default). Aggregate exposures to a single obligor into a single exposure, so each obligor in the portfolio maps to a single exposure. Let subscript i index the obligor (or exposure). Let ζ_i be the initial rating grade of obligor i. For each rating grade, let p_ζ be the associated unconditional probability of default at the risk horizon.

Let x be the vector of CreditMetrics risk factors, e.g., stock market indices, and let Σ be the variance-covariance matrix of x. Without loss of generality, assume there are ones on the diagonal of Σ, so the marginal distributions are all $\mathrm{N}(0,1)$. The condition of obligor i is represented by a latent variable y_i, given by

$$y_i = xw_i + \sigma_i \epsilon_i$$

where w_i is a vector of factor loadings for obligor i, ϵ_i is an idiosyncratic $\mathrm{N}(0,1)$ random variable, and σ_i is a weight on the idiosyncratic effect. Without loss of generality, it is imposed that y_i has variance 1 (i.e., that $w_i'\Sigma w_i + \sigma_i^2 = 1$). Associated with each rating grade is a "cut-off value" C_ζ.[1] When the latent variable falls under the cut-off, the obligor defaults. That is, default occurs if

$$xw_i + \sigma_i \epsilon_i < C_{\zeta(i)}.$$

The C_ζ values are set so that $p_\zeta = \Phi(C_\zeta)$. Let $p_i(x)$ be obligor i's probability of default conditional on a realization of x. This is given by

$$p_i(x) = \Phi((C_{\zeta(i)} - xw_i)/\sigma_i)$$

where Φ is the standard normal cdf.

The CreditRisk$^+$ methodology can now be applied in a straightforward manner. We first derive the conditional probability generating function $F(z|x)$ for the total number of defaults in the portfolio. Conditional on x, default events are independent across obligors. Therefore,

$$\begin{aligned} F(z|x) = \prod_i F_i(z|x) &= \prod_i (1 - p_i(x) + p_i(x)z) \\ &\approx \prod_i \exp(p_i(x)(z-1)) = \exp(\mu(x)(z-1)) \end{aligned} \tag{1}$$

where $\mu(x) \equiv \sum_i p_i(x)$. To get the unconditional probability generating function $F(z)$, we integrate out the x:

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

where ϕ_Σ is the multivariate $\mathrm{N}(0,\Sigma)$ pdf. The unconditional probability that exactly n defaults

[1] In the CreditMetrics Technical Document, the letter Z is used to denote cut-off values.

will occur in the portfolio is given by the coefficient on z^n in the Taylor series expansion of $F(z)$:

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

These integrals are analytically intractable, and in practice would be approximated using Monte Carlo techniques.

The final step in CreditRisk$^+$ is to obtain the probability generating function $G(z)$ for losses.[2] Let L_i be the loss given default of exposure i. Round the values of L_i to the nearest element in the set of "standardized exposure" levels $\{\nu_1, \ldots, \nu_m\}$. Divide the portfolio into subportfolios S_j ($j = 1, \ldots, m$) based on exposure size; i.e., obligor i is in S_j if and only if $L_i = \nu_j$. Let G_j denote the probability generating function for losses within S_j. The probability of a loss of $n\nu_j$ in S_j must equal the probability of n defaults within S_j, so the coefficient on $z^{n\nu_j}$ in the expansion of $G_j(z|x)$ must equal the coefficient on z^n in the expansion of $F_{S_j}(z|x)$:

$$\begin{aligned} G_j(z|x) &= F_{S_j}(z^{\nu_j}|x) = \prod_{i \in S_j} F_i(z^{\nu_j}|x) \\ &= \prod_{i \in S_j} \exp(p_i(x)(z^{\nu_j} - 1)) = \exp(\mu_j(x)(z^{\nu_j} - 1)) \end{aligned}$$

where $\mu_j(x) \equiv \sum_{i \in S_j} p_i(x)$. Because default events are conditionally independent, subportfolio losses are conditionally independent, and thus

$$G(z|x) = \prod_j G_j(z|x) = \exp(\sum_j \mu_j(x)(z^{\nu_j} - 1)) = \exp(\mu(x)(P(z|x) - 1))$$

where $P(z|x) \equiv \left(\sum_j \mu_j(x) z^{\nu_j} \right) / \mu(x)$. Finally, we integrate out x to get the unconditional $G(z)$.

[2] The remaining derivations are drawn without modification from the CreditRisk$^+$ document, appendixes A3.4 and A9.2.

The unconditional probability that there will be n units of loss in the total portfolio is given by the coefficient on z^n in the Taylor series expansion of $G(z)$.

2 Mapping CreditRisk$^+$ to the CreditMetrics framework

Translating in the opposite direction is equally straightforward. Let x be the vector of CreditRisk$^+$ risk factors. As is CreditRisk$^+$, assume the risk factors are orthogonal and that x_k $(k = 1, \ldots, n)$ is distributed Gamma(α_k, β_k). Let $\mu_k = \alpha_k \beta_k$ be the expected value of x_k.[3]

We take the most general representation of the CreditRisk$^+$ model in that we allow the default probability of obligor i to depend on more than one factor. Let w_i be the n-dimensional simplex representing risk-factor loadings for obligor i.[4] The conditional probability of default for obligor i is given by

$$p_i(x) = p_{\zeta(i)} \sum_{k=1}^{n} \frac{x_k}{\mu_k} w_{ik}.$$

As in the single factor CreditRisk$^+$ model, the unconditional default probability is $p_{\zeta(i)}$.

To move into the CreditMetrics framework, we assign to obligor i a latent variable y_i defined by:

$$y_i = \left(\sum_{k=1}^{n} \frac{x_k}{\mu_k} w_{ik} \right)^{-1} \epsilon_i. \tag{2}$$

The idiosyncratic risk factors ϵ_i are independently and identically distributed Exponential with parameter 1. Obligor i defaults if and only if $y_i < p_{\zeta(i)}$. The conditional probability of default is given by:

$$\Pr(y_i < p_{\zeta(i)} | x) \quad = \quad \Pr\left(\epsilon_i < p_{\zeta(i)} \sum_{k=1}^{n} \frac{x_k}{\mu_k} w_{ik} | x \right)$$

[3]See equation (50) in the CreditRisk$^+$ documentation for the density of the gamma distribution.
[4]In the CreditRisk$^+$ documentation, the letter θ is used to denote the factor weights.

$$
\begin{aligned}
&= 1 - \exp\left(-p_{\zeta(i)} \sum_{k=1}^{n} \frac{x_k}{\mu_k} w_{ik}\right) \\
&\approx p_{\zeta(i)} \sum_{k=1}^{n} \frac{x_k}{\mu_k} w_{ik} = p_i(x) \qquad (3)
\end{aligned}
$$

where the second line follows using the cdf for the exponential distribution, and the last line relies on the same approximation formula as equation (1) above. The unconditional probability of default is simply $p_{\zeta(i)}$.

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.[5] In CreditMetrics, the cut-off values C_ζ are determined as functions of the associated unconditional default probabilities p_ζ. Here, the cut-off values are simply the p_ζ. Other than these differences in form, the process is identical. A single portfolio simulation would consist of a single set of random draws of sector risk-factors and a single set of random draws of idiosyncratic risk-factors. From these, the obligors' latent variables are calculated, and these in turn determine default events.

3 Simulation exercises

[TO BE COMPLETED. This section of the paper will present a series of simulation exercises in which the two models are calibrated to the same simulated "test-deck" portfolio. I will show that the two models can deliver the same mean and standard deviation of loss, but differ considerably in the tails. The goal will be to identify the assumptions in each of the models which account for the tail differences.]

[5] One could quasi-linearize equation (2) 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.

Discussion

This paper demonstrates that there is no unbridgeable difference in the views of default 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 only sources of discrepancy in results are differences in distributional assumptions, functional forms, and calibration methods.

The restrictions placed on CreditMetrics are revealing. Obviously, CreditRisk$^+$ can be compared only to a two-state version of CreditMetrics. It is arguable whether a multi-state model is preferable to a two-state model for the hold-to-maturity loan book, but the two-state restriction certainly makes it more difficult to incorporate traded and non-traded positions into the same risk framework. The assumption in CreditRisk$^+$ of fixed exposure size may be even more restrictive. Research in progress investigates whether the CreditMetrics assumption of *idiosyncratic* risk in loss given default is adequate, or whether instead there is a significant correlation between default rates and loss given default. Regardless of the findings of that study, the presence of purely idiosyncratic risk would not be immaterial. Even a large portfolio will have relatively few defaults, so the law of large numbers does not imply that idiosyncratic risk in loss given default is diversified away in a typical large bank portfolio.

It should be noted that there is no loss of generality in the assumption of independence across sector risk-factors in CreditRisk$^+$. In each model, the vector of factor loads (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.[6] However, the need to impose orthogonality in CreditRisk$^+$ does imply that greater care must be given to identifying and calibrating sectoral risks in that

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

model.

Finally, equations (1) and (3) serve to emphasize that CreditRisk$^+$ relies on an approximation formula that holds only for low default probabilities. If the average credit quality of a portfolio is poor, say B3 or worse, the approximation formula performs poorly, and CreditRisk$^+$ may even produce default probabilities greater than one.

References

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

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