Calculation of Higher Moments in CreditRisk$^+$ with Applications

Michael B. Gordy*

Board of Governors of the Federal Reserve System

First draft: June 9, 1999

This draft: September 21, 2001

Abstract

CreditRisk$^+$ is an influential and widely implemented model of portfolio credit risk. As a close variant of models long used for insurance risk, it retains the analytical tractability for which the insurance models were designed. Value-at-risk can be obtained via a recurrence-rule algorithm, so Monte Carlo simulation can be avoided. Little recognized, however, is that the algorithm is fragile. Under empirically realistic conditions, numerical error can accumulate in the execution of the recurrence rule and produce wildly inaccurate results for value-at-risk.

This paper provides new tools for users of CreditRisk$^+$ based on the cumulant generating function ("cgf") of the portfolio loss distribution. Direct solution for the moments of the loss distribution from the cgf is almost instantaneous and is computationally robust. Thus, the moments provide a convenient, quick and independent diagnostic on the implementation and execution of the standard solution algorithm. Better still, with the cgf in hand we have an *alternative* to the standard algorithm. I show how tail percentiles of the loss distribution can be calculated quickly and easily by saddlepoint approximation. On a large and varied sample of simulated test portfolios, I find a natural complementarity between the two algorithms: Saddlepoint approximation is accurate and robust in those situations for which the standard algorithm performs least well, and is less accurate in those situations for which the standard algorithm is fast and reliable.

* The views expressed herein are my own and do not necessarily reflect those of the Board of Governors or its staff. I thank Tom Wilde and Paul Embrechts for helpful guidance and Gary Anderson for generous assistance with *Mathematica*.

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⟩.

The CreditRisk$^+$ model of portfolio credit risk has drawn significant practitioner interest since its publication in 1997 by Credit Suisse Financial Products.[1] It appears to be currently the most widely-implemented of the so-called *two-state* or *default-mode* models, in which loss is recognized only in the event of borrower default before a fixed horizon date (typically one year ahead). Loss due to change in market value is ignored in these models. In particular, changes in borrower credit quality short of default, such as might be connoted by changes in agency rating, have no impact. The trend among practitioners at larger institutions appears to be towards adoption of *multi-state* models in which loss is defined on a mark-to-market basis. Multi-state modeling is indeed essential for pricing and trading applications and for proper treatment of the effect of loan maturity on credit risk. Nonetheless, two-state models do have advantages in the context of a buy-and-hold lending portfolio. The default-mode definition of loss is more consistent with traditional book-value accounting, upon which the legal and regulatory standards of bank solvency depend. Two-state models may also be easier to calibrate using data readily available within the bank. Moreover, lending contracts often embed options that may be difficult to value precisely, especially in the context of a long-term and multi-faceted banking relationship.[2] Pricing rules in the current generation of mark-to-market credit value-at-risk models are relatively simple and thus may provide a spurious sense of precision for some portfolios. As a practical matter, unless the portfolio is of very high quality or of long average duration, default risk is the dominant component of credit risk, so a default-mode measure of value-at-risk may be an adequate characterization of credit risk for many purposes.[3]

As a close variant of models long used for insurance risk, CreditRisk$^+$ retains the analytical tractability for which the insurance models were designed. It yields a closed form solution for the probability generating function ("pgf") of the distribution of portfolio credit loss. From this pgf, the probabilities associated with each possible level of credit loss can be calculated successively via a recurrence relationship. Provided that one chooses a reasonably coarse discretization of exposure sizes and parsimonious set of risk factors, the calculations can be performed quickly and accurately.

This paper provides new tools for users of CreditRisk$^+$ based on the cumulant generating function of the portfolio loss distribution. I show that the CreditRisk$^+$ cumulant generating function ("cgf") can be calculated quite easily and poses no numerical challenges. From the cgf, one can directly obtain the moments of the loss distribution. Though solution of tail percentiles in CreditRisk$^+$ can be fast, direct solution of the moments can be orders of magnitude faster. Whereas solution time for tail percentiles increases exponentially with the number of risk factors and as a polynomial of the discretization of exposure sizes, solution time for moment calculations is invari-

[1]See Credit Suisse Financial Products (1997) for a complete description and technical discussion.

[2]For example, loan covenants may appear to give the bank power to renegotiate from a position of strength if, say, borrower operating income deteriorates, yet are routinely waived in practice.

[3]Loans to highly-rated borrowers typically have very small risk of default within the model horizon, so most of the credit risk in long-maturity, high-grade loans is associated with downgrade risk rather than default risk. In this case, default-mode measures may be seriously misleading. Lending books at most commercial banks are dominated by short to medium-term lending to borrowers of moderate credit quality.

ant with respect to discretization of exposure size and linear in the number of risk factors.[4] More importantly, the numerical accuracy of the moment calculations is unaffected by the number of risk factors and is robust to arbitrarily fine discretization of exposure sizes. Therefore, the moments can be used as a convenient, quick and independent diagnostic on the implementation and execution of the standard CreditRisk$^+$ solution algorithm.

Better still, with the cgf in hand we have an *alternative* to the standard recurrence-rule algorithm. I show how tail percentiles of the loss distribution can be calculated quickly and easily by saddlepoint approximation. On a large and varied sample of simulated test portfolios, I find a natural complementarity between the two algorithms: Saddlepoint approximation is accurate and robust in those situations for which the standard algorithm performs least well, and is less accurate in those situations for which the standard algorithm is fast and reliable.

The CreditRisk$^+$ model is summarized in Section 1. In Section 2, I derive the CreditRisk$^+$ cumulant generating function and show how to obtain from it the mean, variance, and indexes of skewness and kurtosis. Section 3 explains why numerical issues can arise in the standard algorithm for calculating the CreditRisk$^+$ loss distribution, and shows how the moments provide a useful diagnostic. Saddlepoint approximation of tail percentile values of the loss distribution is described and tested in Section 4.

1 Summary of model and notation

Portfolio credit risk in CreditRisk$^+$ is assumed to be driven by a vector X of "risk factors" that represent the random component of macroeconomic conditions from the current time through the model horizon (typically one year). Conditional on a particular realization $X = x$, defaults of individual obligors are assumed to be independently distributed Bernoulli draws. In the simplest one-factor version of the model, the probability of drawing a default for obligor i conditional on $X = x$ is given by $p_i(x) = \bar{p}_i x$, where $\bar{p}_i$ is the unconditional default probability of obligor i (for example, as indicated by its rating agency grade). The single risk factor, which represents uncertainty in the broad credit cycle, is thus assumed to have a proportional effect on the default probability of every obligor. When a bad outcome is realized ($X > 1$), default probabilities are scaled up together, so the portfolio is likely to experience higher than expected levels of loss. When a good outcome is realized ($X < 1$), default probabilities are scaled down, and portfolio loss is likely to remain below its unconditional expected value.

Though not infrequently implemented in this simple form, the proportional single-factor specification imposes a highly constrained structure on default correlations. As conditions do not move in lockstep across sectors of the economy, multiple risk-factors are needed to capture industry and regional effects. Furthermore, the strict proportionality assumption should be relaxed. Calibration exercises such as in Gordy (2000) suggest that high-rated obligors tend to have greater sensitivity

[4] To be precise, if the "standardized loss unit" is halved, solution time for the full model is increased by 2^{K+1}, where K is the number of risk factors.

to systematic risk than low-rated obligors. Therefore, it is increasingly common to work with the generalized K-factor specification

$$p_i(x) = \bar{p}_i \left(w_{i0} + \sum_{k=1}^{K} x_k w_{ik} \right). \tag{1}$$

The X_k are assumed to be independent gamma variates with mean one and variance $1/\tau_k$.[5] The vector of "factor loadings" $(w_{i1}, \ldots, w_{iK})$ measures the sensitivity of obligor i to each of the risk factors. The constant w_{i0} is a loading on idiosyncratic risk. Because obligor-specific risks are diversifiable, they jointly contribute negligible volatility to a well-diversified portfolio. Therefore, the CreditRisk$^+$ manual (§A12.3) suggests that w_{i0} be construed as a weighting on a degenerate "specific risk" sector with $X_0 = 1$ and $\tau_0 = \infty$. The weights w_{ik} $(k = 0, \ldots, K)$ are required to sum to one for each obligor, which guarantees that $\mathrm{E}[p_i(X)] = \bar{p}_i$.

The distribution of losses in CreditRisk$^+$ is characterized by its probability generating function. The pgf $\mathcal{F}_y(z)$ of a discrete random variable Y is a function of an auxilliary variable z such that the probability that $Y = n$ is given by the coefficient on z^n in the polynomial expansion of $\mathcal{F}_y(z)$. Conditional on $X = x$, default for obligor i is a Bernoulli random variable with parameter $p_i(x)$, and so has pgf

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

where the approximation is valid as long as $p_i(x)$ is fairly small.[6]

In the event of default by obligor i, there is a fixed loss of ν_i, which is constrained to be an integer multiple of a "standardized loss unit," e.g., one million dollars.[7] The conditional pgf for loss on obligor i is then simply $\mathcal{G}_i(z|x) = \mathcal{F}_i(z^{\nu(i)}|x)$. The conditional pgf of a sum of independent random variables is the product of the individual conditional pgfs, so the conditional pgf for portfolio loss

[5] This is a variant on the presentation in the CreditRisk$^+$ manual, in which X_k has mean μ_k and the conditional probabilities are given by $p_i(x) = \bar{p}_i(\sum w_{ik}(x_k/\mu_k))$. In my presentation, which follows Gordy (2000), the constants $1/\mu_k$ are absorbed into the normalized X_k without any loss of generality.

[6] The last equality is called 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 analytical facility of the model.

[7] Bürgisser, Kurth and Wagner (2001) extend the standard model to introduce idiosyncratic and systematic risk in loss severity.

is

$$\begin{aligned}
\mathcal{G}(z|x) &= \prod_i \mathcal{G}_i(z|x) \approx \exp(\sum_i p_i(x)(z^{\nu(i)}-1)) = \exp\left(\sum_i \bar{p}_i\left(w_{i0}+\sum_k w_{ik}x_k\right)(z^{\nu(i)}-1)\right) \\
&= \exp\left(\sum_i \bar{p}_i w_{i0}(z^{\nu(i)}-1))\right)\exp\left(\sum_{k=1}^{K} x_k(\sum_i \bar{p}_i w_{ik}(z^{\nu(i)}-1))\right) \\
&= \exp(\mu_0(\mathcal{P}_0(z)-1))\exp\left(\sum_{k=1}^{K} x_k\mu_k(\mathcal{P}_k(z)-1)\right) \qquad (3)
\end{aligned}$$

where $\mu_k \equiv \sum_i w_{ik}\bar{p}_i$ and

$$\mathcal{P}_k(z) \equiv \frac{1}{\mu_k}\sum_i w_{ik}\bar{p}_i z^{\nu(i)}. \qquad (4)$$

By integrating over the x_k, we obtain the unconditional pgf:

$$\mathcal{G}(z) = \exp(\mu_0(\mathcal{P}_0(z)-1))\prod_{k=1}^{K}\left(\frac{1-\delta_k}{1-\delta_k\mathcal{P}_k(z)}\right)^{\tau_k} \qquad (5)$$

where $\delta_k \equiv \mu_k/(\tau_k+\mu_k)$.

The unconditional probability that there will be ℓ standardized units of loss in the total portfolio is given by the coefficient on z^ℓ in the Taylor series expansion of $\mathcal{G}(z)$. The CreditRisk$^+$ manual (§A.10) describes how the coefficients may be calculated via a simple recurrence relation.

2 The CreditRisk$^+$ cumulant generating function

The cumulant generating function ("cgf") of a distribution is the log of its moment generating function, which in turn can be expressed in terms of the probability generating function. If $\mathcal{F}_y(z)$ is the pgf of Y, then the cgf $\psi_y(z)$ is given by

$$\psi_y(z) = \log(\mathcal{F}_y(\exp(z))). \qquad (6)$$

The jth cumulant of Y is given by the jth derivative of ψ_y evaluated at $z=0$. Cumulants are related to the more familiar raw moments and central moments, but for many applications have more useful properties.[8] For the first four cumulants of a distribution, the relationship to the central moments is straightforward. Let m be the mean of the distribution for some random variable Y, and let m_j be the jth centered moment of Y, i.e., $m_j = \mathrm{E}[(Y-m)^j]$. The first four cumulants are given by $\kappa_1 = m$, $\kappa_2 = m_2$, $\kappa_3 = m_3$, and $\kappa_4 = m_4 - 3m_2^2$. Thus, the first cumulant is the mean

[8] See Stuart and Ord (1994, §3.12) for more information on cumulants.

and the second is the variance. The index of skewness and index of kurtosis can be written

$$\text{Skewness}(Y) \equiv \frac{m_3}{m_2^{3/2}} = \frac{\kappa_3}{\kappa_2^{3/2}}, \qquad \text{Kurtosis}(Y) \equiv \frac{m_4}{m_2^2} = \frac{\kappa_4}{\kappa_2^2} + 3.$$

After the first four terms, the relationship becomes increasingly complex; e.g., $\kappa_5 = m_5 - 10m_2m_3$.

For the CreditRisk$^+$ loss distribution, we find

$$\psi(z) = \log(\mathcal{G}(\exp(z))) = \mu_0(\mathcal{P}_0(z) - 1) + \sum_{k=1}^{K} \tau_k \cdot \log\left(\frac{1-\delta_k}{1-\delta_k \mathcal{P}_k(\exp(z))}\right). \tag{7}$$

Substituting the function

$$\mathcal{Q}_k(z) \equiv \mu_k \mathcal{P}_k(\exp(z)) = \sum_i w_{ik}\bar{p}_i \exp(\nu_i z) \tag{8}$$

into equation (7), we get

$$\psi(z) = (\mathcal{Q}_0(z) - \mu_0) + \sum_{k=1}^{K} \tau_k \cdot \log\left(\frac{\mu_k(1-\delta_k)}{\mu_k - \delta_k \mathcal{Q}_k(z)}\right) \equiv \psi_0(z) + \sum_{k=1}^{K} \psi_k(z) \tag{9}$$

where the functions $\psi_k(z)$ are introduced for notational convenience.

Let D denote the differential operator (i.e., $D^j f(x)$ is the jth derivative of f with respect to x). For the specific risk sector, the j^{th} derivative of $\psi_0(z)$ is simply the j^{th} derivative of $\mathcal{Q}_0(z)$, which can be evaluated using the general equation

$$D^j \mathcal{Q}_k(z) = \sum_i w_{ik}\bar{p}_i \nu_i^j \exp(\nu_i z) \quad \forall j \geq 0, k \in \{0, 1, \ldots, K\}. \tag{10}$$

For the systematic risk sectors, the first derivative of ψ_k is

$$\psi'_k(z) = \tau_k \left(\frac{\delta_k D\mathcal{Q}_k(z)}{\mu_k - \delta_k \mathcal{Q}_k(z)}\right). \tag{11}$$

It is useful to generalize the expression in parenthesis as

$$V_{j,k}(z) \equiv \frac{\delta_k D^j \mathcal{Q}_k(z)}{\mu_k - \delta_k \mathcal{Q}_k(z)} \tag{12}$$

and to observe that derivatives of V have the simple recurrence relation

$$DV_{j,k}(z) = V_{j+1,k}(z) + V_{j,k}(z)V_{1,k}(z). \tag{13}$$

Therefore, using

$$\psi_k'(z) = \tau_k V_{1,k}(z) \tag{14a}$$

and equation (13), higher derivatives of ψ_k can be generated mechanically. The second, third and fourth are given by

$$\psi_k''(z) = \tau_k(V_{2,k}(z) + V_{1,k}(z)^2) \tag{14b}$$
$$\psi_k'''(z) = \tau_k(V_{3,k}(z) + 3V_{2,k}(z)V_{1,k}(z) + 2V_{1,k}(z)^3) \tag{14c}$$
$$\psi_k''''(z) = \tau_k(V_{4,k}(z) + 4V_{3,k}(z)V_{1,k}(z) + 3V_{2,k}(z)^2 + 12V_{2,k}(z)V_{1,k}(z)^2 + 6V_{1,k}(z)^4). \tag{14d}$$

Appendix A shows how to automate calculation of higher derivatives of ψ_k using *Mathematica*. The $V_{j,k}(z)$ are easily evaluated using equations (8) and (10).

Regardless of portfolio composition, the cgf always satisfies certain properties. So long as all the w_{ik}, $\bar{p}_i$ and ν_i are nonnegative and there exist loans for which the product $w_{ik}\bar{p}_i\nu_i$ is strictly positive, then $\mathcal{Q}_k$ and its derivatives are positive, continuous, increasing, convex functions of z. Let z_k^* be the unique solution to $\mathcal{Q}_k(z) = \mu_k + \tau_k$ and let $z^* \equiv \min\{z_1^*, \ldots, z_K^*\}$. For all $j \geq 1$, the denominator in $V_{j,k}(z)$ is positive and decreasing for all $z < z_k^*$ and equal to zero at z_k^*. The numerator in $V_{j,k}(z)$ is positive and increasing, so $V_{j,k}(z)$ is positive and increasing for all $z \leq z_k^*$. Because $\mathcal{Q}_k(0) = \mu_k$, we must have $z_k^* > 0$ for all $k = \{1, \ldots, K\}$. It follows that $\psi'(z)$ and higher derivatives are positive, continous, increasing and convex for all $z < z^*$ and that the valid domain of ψ contains the origin. As $z \to z^*$, $\psi(z)$ and all its derivatives tend to ∞.

To obtain the cumulants, I evaluate the derivatives of ψ at $z = 0$. Using the equality $\mathcal{Q}_k(0) = \mu_k$ and the definition of δ_k, I get

$$V_{j,k}(0) = (1/\tau_k)D^j\mathcal{Q}_k(0), \tag{15}$$

and substitute into equations (14a)–(14d).

Calculation of the cgf and the cumulants is extremely fast and accurate. Though it may appear somewhat tedious, the expressions for $V_{j,k}(z)$ are trivially simple to program. Calculating the first four cumulants for a portfolio of thousands of obligors takes only a small fraction of a second.

3 Cumulants as a diagnostic tool

Direct calculation of cumulants can provide users and outside auditors a fast and robust tool for checking the software implementation and execution of CreditRisk$^+$. Although implementation errors ought to occur only during the development stage, execution errors may arise at any time. For example, too coarse a discretization of exposure sizes causes a loss of accuracy in model results. One can use the cumulants to check quickly the sensitivity of the moments of the loss distribution to the choice of standardized loss unit.

A more likely and more dangerous source of execution error in CreditRisk$^+$ is numerical roundoff error. The standard solution technique, known as Panjer recursion in the insurance literature, requires that we express the log derivative of $\mathcal{G}(z)$ as a rational polynomial. That is, we first solve for polynomials $A(z)$ and $B(z)$ such that

$$\frac{A(z)}{B(z)} = \frac{d}{dz}(\log \mathcal{G}(z)) = \mu_0 \mathcal{P}_0'(z) + \sum_{k=1}^{K} \frac{\delta_k \mathcal{P}_k'(z)}{1 - \delta_k \mathcal{P}_k(z)}. \tag{16}$$

If $\mathcal{P}_k(z)$ is of degree m_k, then the polynomial $B(z)$ is of degree

$$s = m_1 m_2 \cdots m_K$$

and $A(z)$ is of degree

$$r = s\left((m_0 - 1) + \frac{m_1 - 1}{m_1} + \ldots \frac{m_K - 1}{m_K}\right) \approx s(K + m_0 - 1)$$

when the m_k are large. The degree m_k is determined by the largest loan with positive weighting on risk factor k, that is, $m_k = \max\{\nu_i : w_{ik} > 0\}$. Therefore, the finer the discretization of loss exposure sizes, the larger are the m_k, and the longer are the polynomials $A(z)$ and $B(z)$. In real applications, it should not be unusual to have m_k equal to 100 or 200, so r and s can become enormous with only a few systematic risk sectors. Except in the special case of a single systematic risk sector and no specific risk sector, both $A(z)$ and $B(z)$ contain coefficients of both signs. As roundoff errors occur when summing numbers of similar magnitude but opposite sign, numerical accuracy may be lost in the polynomial convolutions needed to simplify the sum of rational polynomials on the right-hand-side of equation (16) into the single rational polynomial $A(z)/B(z)$.[9]

As demonstrated in the CreditRisk$^+$ manual, the expansion of $\mathcal{G}(z)$ as $g_0 + g_1 z + g_2 z^2 + \ldots$ follows the recurrence relation

$$g_{j+1} = \frac{1}{b_0(j+1)}\left(\sum_{i=0}^{\min(r,j)} a_i g_{j-i} \; + \sum_{i=0}^{\min(s,j)-1} b_{i+1}(j-i) g_{j-i}\right) \tag{17}$$

where the $\{a_i\}$ and $\{b_i\}$ are the coefficients of $A(z)$ and $B(z)$ respectively. This rule is iterated until we reach the number of loss units j^* such that $G(j^*) = g_0 + g_1 + \ldots + g_{j^*}$ covers the targeted bank solvency probability. As solvency targets are typically on the order of 99.8% to 99.97%, there may be several thousand iterations before halting. With each iteration, roundoff error accumulates. At the far tail of the distribution, where the g_j are small and declining slowly, numerical issues may become especially significant; indeed, calculated probabilities at the far tail can be negative. The potential for error increases as the skewness of the loss exposure distribution is increased, as

[9] Rounding error induced by subtractive cancellation is perhaps the most common source of numerical error in computation. See Press, Teukolsky, Vetterling and Flannery (1992, §1.3).

the discretization of exposure sizes is made finer and as the number of risk-factors is increased. Fortunately, these numerical issues do not arise in direct calculation of cumulants, because all the summations used in calculating the cumulants contain only nonnegative terms of roughly similar magnitude. These calculations are fast and accurate, regardless of model parameters and portfolio make-up.

Implementation of a diagnostic based on the cumulants is straightforward. I assume that the target quantile q for value-at-risk is sufficiently far into the tail that extreme value theory can be applied. As shown in Appendix C, the tail of the CreditRisk$^+$ loss distribution can be approximated by the tail of an exponential distribution with parameter β. By construction, $G(j^*) \approx q$, so

$$q \approx G(j^*) \approx 1 - \exp(-j^*/\beta)$$

which implies that a consistent estimator of β (as we move farther into the tail) is

$$\beta = -j^*/\log(1-q).$$

For the exponential distribution,

$$\mathrm{E}[Z^r | Z > u] = \sum_{m=0}^{r} \binom{r}{m} m! \beta^m u^{r-m}$$

which implies that

$$\mathrm{E}[Y^r | Y > j^*] = \sum_{m=0}^{r} \binom{r}{m} m! (-j^*/\log(1-q))^m j^{*r-m} = j^{*r} \sum_{m=0}^{r} \binom{r}{m} \frac{m!}{(-\log(1-q))^m}.$$

Therefore, we can approximate

$$\begin{aligned} \mathrm{E}[Y^r] &= \mathrm{E}[Y^r | Y \leq j^*] \Pr(Y \leq j^*) + \mathrm{E}[Y^r | Y > j^*] \Pr(Y > j^*) \\ &\approx \sum_{j \leq j^*} g_j j^r + (1-q) j^{*r} \sum_{m=0}^{r} \binom{r}{m} \frac{m!}{(-\log(1-q))^m}. \end{aligned} \tag{18}$$

The first term is produced as a "by-product" of the Panjer recursion. As probabilities g_j are calculated via recurrence for successive values of j, the values $g_j j^r$ are cumulated. The second term represents the correction for values of $j > j^*$.

If the Panjer recursion is executed without significant numerical error, then the moments calculated in equation (18) should match the uncentered moments calculated from the cumulants, so the absolute relative error for each order r should be negligible. As a rule of thumb, I find that a mismatch of over 1% on the mean or 2.5% on the standard deviation strongly indicates that the Panjer algorithm has failed.[10]

[10] For this purpose, it is best to calculate cumulants using the same discretized loss exposure values as used in the Panjer recursion. If cumulants are calculated using the exact loss exposures, then the diagnostic could fail due to

4 Saddlepoint approximation of value-at-risk

Edgeworth expansions are frequently used to approximate distributions that lack a convenient closed-form solution but for which moments are available. It is widely recognized that this classic technique works well in the center of a distribution, but can perform very badly in the tails. Indeed, it often produces negative values for densities in the tails. Saddlepoint expansion can be understood as a refinement of Edgeworth expansion. While not as widely applied, it offers vastly improved performance in every respect. In saddlepoint approximation, the density of a distribution $f(x)$ at each point x is obtained by "tilting" the distribution around x to get a new distribution $\tilde{f}(z;x)$ such that x is in the center of the new distribution. The tilted distribution is approximated via Edgeworth expansion, and the mapping is inverted to obtain the approximation to $f(x)$. Because Edgeworth expansion is used only at the center of the tilted distributions, the approximation is in general uniformly good, in the sense of having a small relative error even in the tails. A consequence is that saddlepoint approximation of the cumulative distribution function is often quite good throughout the tail. Indeed, for some distributions, including the gamma, the approximation is in fact exact. Jensen (1995) provides an excellent textbook treatment of the subject.

Although saddlepoint approximation appeared first in the 1930s and was fully developed by the mid-1980s, it is not widely used because it requires that the cumulant generating function have tractable form. It is not often the case that a distribution of interest with intractable cdf has a known and tractable cgf. Nonetheless, there have been efforts to apply saddlepoint approximation to problems in risk management. Feuerverger and Wong (2000) derive a cgf of the distribution of a delta-gamma approximation to loss in a value-at-risk model of market risk, and apply saddlepoint approximation using this approximated cgf. Martin, Thompson and Browne (2001) provide a simple model of credit risk that is suitable for saddlepoint approximation. However, their model is designed around the need to maintain tractability rather than ease and reliability of calibration. In this section, I show that saddlepoint approximation applies quite well to CreditRisk$^+$. Accurate approximations to the tail percentiles of the loss distribution can be obtained with trivial computation time.

Different expansions lead to different forms for the saddlepoint approximation. An especially parsimonious and easily implemented variant is the Lugannani-Rice approximation. Let Y be a random variable with distribution $G(y)$ and cgf $\psi(z)$, and let $\hat{z}$ denote the unique real root of the equation $y = \psi'(z)$. The Lugannani-Rice formula for tail of G is

$$1 - G(y) \approx 1 - \Phi(w) + \phi(w)\left(\frac{1}{u} - \frac{1}{w}\right) \tag{19}$$

where

$$w = \sqrt{2(\hat{z}y - \psi(\hat{z}))} \quad \text{and} \quad u = \hat{z}\sqrt{\psi''(\hat{z})}$$

error induced by discretization rather than numerical error in the solution algorithm.

and where Φ and ϕ denote the cdf and pdf of the standard normal distribution.

In risk-management applications, we are rarely interested in the value of $G(y)$ for particular values of y. Value-at-risk ("VaR") is the inverse problem, i.e., finding the amount of loss y such that $G(y) = q$, where q is the target solvency probability.[11] This is convenient from a computational point of view, as we avoid root-solving for $\hat{z}$. Indeed, the only potential challenge in saddlepoint approximation of CreditRisk$^+$ is in finding the upper bound z^* of the valid domain of ψ. I demonstrate in Appendix B that, as long as $q > G(\text{E}[Y])$, the bound z_k^* for each systematic sector k must satisfy

$$0 < z_k^* < -\log(\delta_k) \Big/ \sum_i \frac{w_{ik}\bar{p}_i}{\mu_k}\nu_i.$$

As $\mathcal{Q}_k(z)$ is strictly increasing in z, finding z_k^* in this bounded interval is computationally trivial. The remainder of the calculations proceed as follows:

1. Set $z^* \equiv \min\{z_1^*, \ldots, z_K^*\}$. Form a fine grid of $\hat{z}$ values in the open interval $(0, z^*)$.

2. At each point in the grid, calculate the corresponding y, w and u from the derivatives of ψ.

3. Form a table of pairs $\langle 1 - G(y), y \rangle$ using the Lugannani-Rice formula. Interpolate to find the value of y corresponding to $1 - G(y) = 1 - q$.

The saddlepoint approximation to VaR always exists. As $\hat{z} \to z^*$, w, u and y tend to ∞. Therefore, for any $q \in (G(\text{E}[Y]), 1)$, we can always find a value of $\hat{z} \in (0, z^*)$ such that the Lugannani-Rice formula for $1 - G(y)$ equals $1 - q$.

Results for a typical portfolio are shown in Figure 1. There are 1000 obligors in this portfolio with an average default probability of 1%. Borrowers are heterogeneous in credit quality, and the loan size distribution is skewed so that the largest 10% of loans account for 40% of total exposure. There are three systematic risk factors, each with $\tau_k = 0.5$. The upper 2% tail of the loss distribution is plotted as a solid line, and the saddlepoint approximation as a dashed line. Approximation error is negligible throughout the tail.

I assess the accuracy and robustness of the saddlepoint approximation on a wide variety of simulated portfolios. The portfolios vary in number of borrowers, credit quality of borrowers, and skewness of exposure size distribution. I also vary the number of systematic sectors and the variance of the systematic risk factors. For each of 15000 test runs:

- I choose the number of obligors n from $\{200, 300, 500, 1000, 5000\}$.

- The average credit quality is specified by choosing a mean default probability θ from the set $\{0.001, 0.005, 0.01, 0.03, 0.05\}$. To ensure heterogeneity within the portfolio, the individual $\bar{p}_i$ are drawn as iid exponential random variables with mean θ.

[11] In the CreditRisk$^+$ model, G is discrete, so technically we are looking for the minimum value of y such that $G(y) \geq q$.

Figure 1: Approximation of the Tail of the Loss Distribution

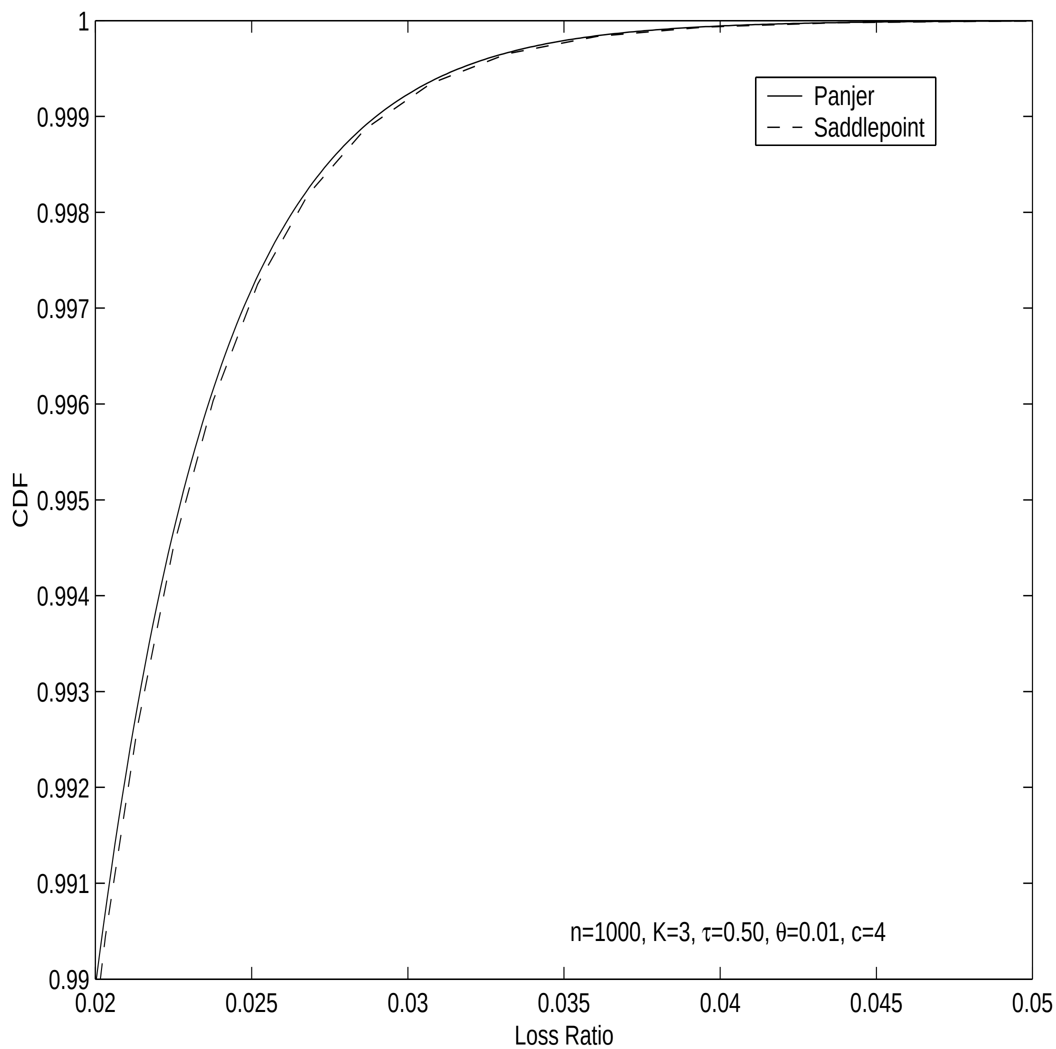

- I choose a "concentration parameter" $c \in \{2, 4\}$. The nominal exposure size of loan i is set to i^c, so the higher value of c creates a more skew distribution of loan sizes. Loss given default is assumed to be 50% of nominal exposure, so loss exposure is $\nu_i = 0.5i^c$.

- The number of systematic sectors K is chosen from $\{1, 3, 7\}$. In order that VaR be roughly invariant with respect to the chosen number of sectors, the variances $1/\tau_k$ must be set larger when K is larger. I draw a uniform random number U and set $1/\tau_k = K(1 + 3U)^2/4$ for all k.

- For each obligor, a vector of factor loadings is drawn as a random simplex on $\Re^{K+1}$.

The model and portfolio are now fully specified. I solve for VaR at the 99.5% percentile via the standard Panjer recursion algorithm. Where that algorithm fails the diagnostic of Section 3, I also estimate VaR via Monte Carlo simulation with 50000 trials. Lastly, I estimate VaR via Lugannani-Rice saddlepoint approximation.

Performance of both the Panjer and saddlepoint methodologies depends strongly on the nature of the portfolio and the complexity of model specification. I found that Panjer recursion almost always passed the moment-based diagnostics *except* when $K = 7$ and $n = 5000$. That is, for simple risk-factor specifications and small to medium sized portfolios, the standard CreditRisk$^+$ solution algorithm appears to work reliably. However, when $K = 7$ and $n = 5000$, the diagnostic failed in 222 out of 1000 simulations. As a practical matter, the consequences for VaR estimation were usually drastic. Compared against the Monte Carlo estimate of VaR, reported VaR from the Panjer algorithm was off by 30% or more in 200 of the 222 detected failures and by 60% or more in 141 of the failures. This result is not unexpected. The greater the number of sectors and the larger the portfolio, the longer the polynomials in the recurrence equation (17), so the greater the opportunity for round-off error to accumulate.[12]

The accuracy of the saddlepoint approximation over my sample of test portfolios is summarized in Table 1. I divide the sample of portfolios by number of systematic sectors and number of obligors. Within each group, I report quantiles of the distribution of absolute relative error ("ARE"). In the first row, for example, we see that the median ARE when $K = 1$ and $n < 400$ is under 0.9%. For 99% of the test portfolios within this group, ARE is 9.22% or less. Maximum ARE in this group is 13.73%. For the case of $n = 5000$ and $K > 1$, I divide the test portfolios into two groups: those for which the Panjer recursion satisfied the moment diagnostic, and those for which I needed to estimate "true" VaR via Monte Carlo.

The table suggests several qualitative conclusions. First, in all cases, saddlepoint approximation is typically quite accurate. For the sample as a whole, median ARE is 0.67%. Indeed, in the typical case, it may well be that the saddlepoint approximation is *more* accurate than Panjer recursion

[12] The number of obligors in not the problem *per se*, but rather the maximum loss exposure among the ν_i. In my test portfolios, higher n is associated with a larger ratio between largest and smallest loans. If I maintain a reasonably fine discretization of the loss exposures, the maximum ν_i becomes quite large. If I choose a more coarse discretization, the algorithm may perform well but the accuracy of the results becomes questionable.

Table 1: Distribution of Absolute Relative Error*

	50%	95%	98%	99%	max
$K = 1$ & $n \in \{200, 300\}$	0.89	5.37	7.59	9.22	13.73
$K = 1$ & $n \in \{500, 1000, 5000\}$	0.80	1.50	1.70	1.91	6.48
$K \in \{3, 7\}$ & $n \in \{200, 300\}$	0.75	7.97	10.87	13.58	22.85
$K \in \{3, 7\}$ & $n = 500$	0.50	3.58	5.78	7.40	18.07
$K \in \{3, 7\}$ & $n = 1000$	0.44	1.84	3.22	4.45	12.84
$K \in \{3, 7\}$ & $n = 5000$ (Panjer)	0.61	1.25	1.54	2.66	5.63
$K \in \{3, 7\}$ & $n = 5000$ (Monte Carlo)	1.17	3.40	4.14	4.63	5.02

*: "True" VaR calculated by Panjer recursion except for final row. Absolute relative error expressed as a percentage of "true" VaR.

because the saddlepoint approximation allows us to use exact values for the loss exposures, rather than discretized values.

Second, saddlepoint approximation is much more reliable for medium to large portfolios than for small ones. For the $K > 1$ test portfolios, the 95% quantile of the distribution of ARE falls from 8% when $n \in \{200, 300\}$ to 3.6% when $n = 500$ to 1.8% when $n = 1000$ to 1.25% when $n = 5000$ (ignoring the cases for which Panjer recursion failed). Higher quantiles of the ARE distributions behave similarly. By $n = 1000$, ARE is under 5% for over 99% of the test portfolios.

Finally, comparing the bottom two rows of the table, ARE is higher when VaR is estimated via Monte Carlo. This should likely be attributed to simulation noise in estimating the "true" VaR rather than to poor performance of the saddlepoint approximation. As the Panjer algorithm was wildly inaccurate for this subsample of test portfolios, we lack an exact solution for comparison.

Results are presented graphically in Figures 2 and 3. In the upper left panel of Figure 2, I plot VaR, as calculated by Panjer recursion, against the saddlepoint approximation for all test portfolios in which $K = 1$. To ease comparison across portfolios of very different sizes, I express VaR as a percentage of total nominal exposure. Small portfolios ($n < 400$) are plotted with a "+" symbol, and large portfolios ($n > 400$) with a "o" symbol. All points lie in a tight band along the $y = x$ line. In the corresponding panel of Figure 3, I plot the ARE values associated with these same test portfolios. Aside from a small cluster associated with small high quality portfolios, ARE is shown to be nearly always bounded under 1%.

In the upper right panels of the figures, I plot the results for $K \in \{3, 7\}$ and $n \in \{200, 300\}$. We see that the saddlepoint approximation performs well in the majority of cases, but that significant discrepancies are not infrequent. Reliability is noticeably better for $n \in \{500, 1000\}$, plotted in the lower left panels. Most of the remaining significant discrepancies are associated with $n = 500$. For the $n = 5000$ test portfolios, plotted in the lower right panels, significant discrepancies are few and nearly always associated with Monte Carlo estimation of VaR (plotted with a "o" symbol).

There are alternatives to the Lugannani-Rice formula. The Barndorff-Nielsen formula (see Feuerverger and Wong 2000) appears to be somewhat less robust for CreditRisk$^+$. Other saddlepoint approximations make use of higher order derivatives of the cgf (e.g., Jensen 1995, §6.1) or

Figure 2: Saddlepoint Approximation of CreditRisk$^+$

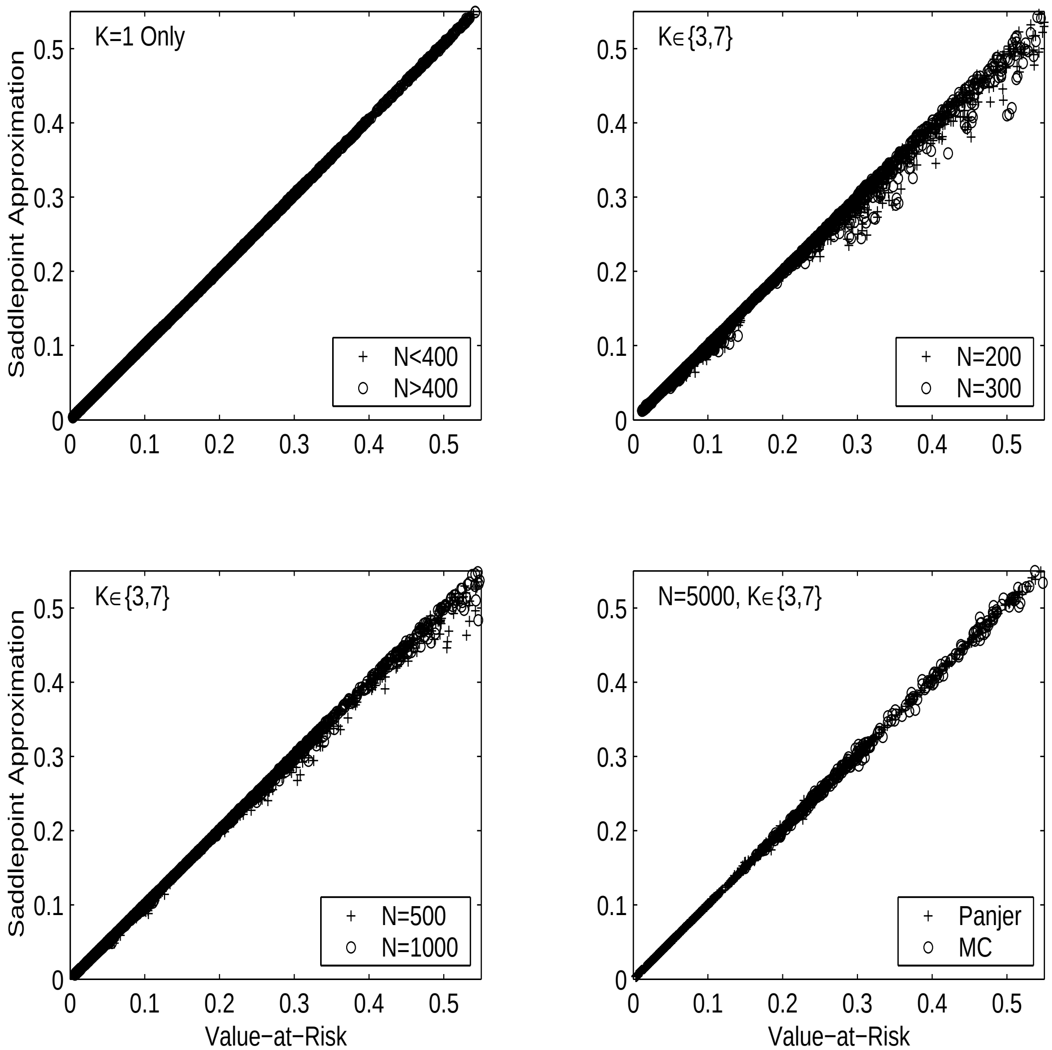

Figure 3: Absolute Relative Error

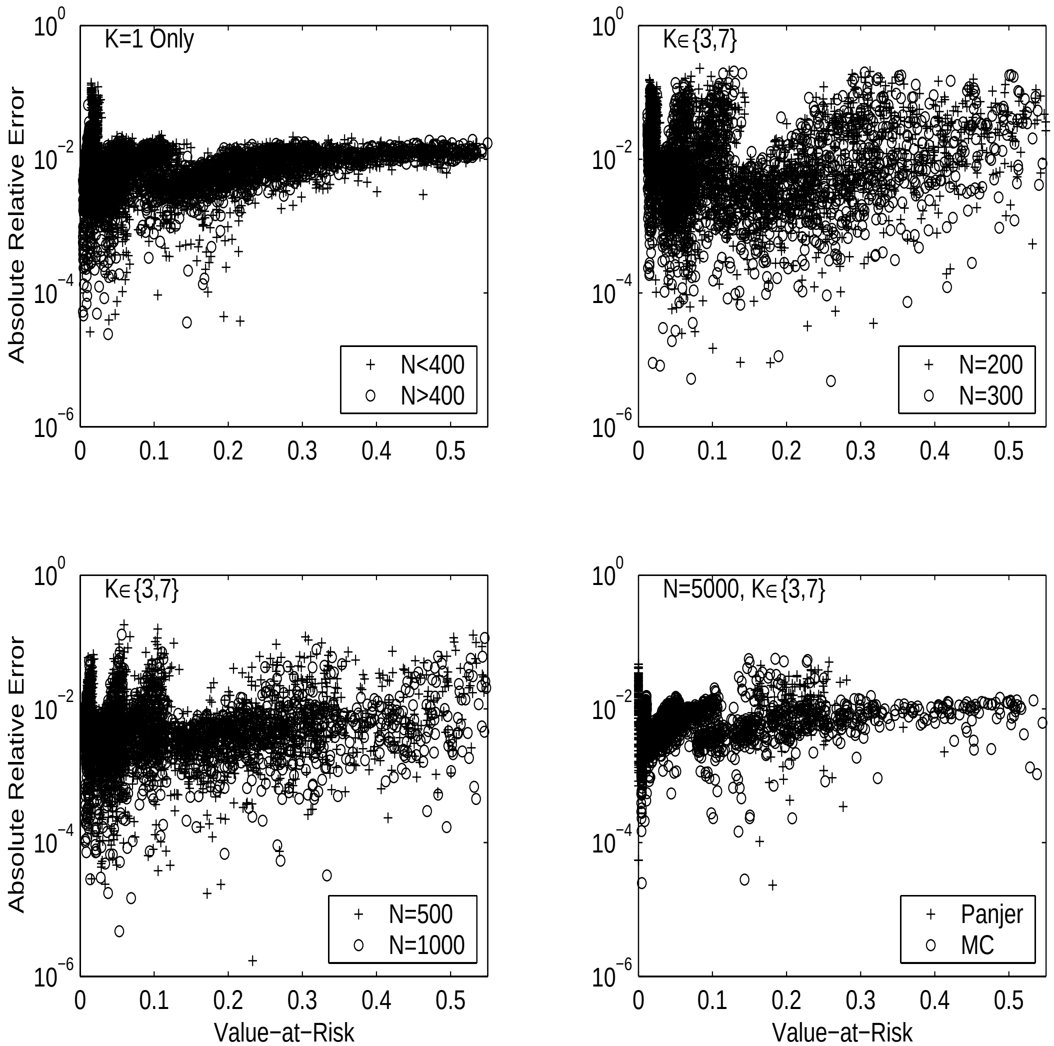

non-Gaussian base distributions (e.g., Wood, Booth and Butler 1993), but do not improve performance in our application.

Conclusion

Compared to the best-known competing models of credit risk, CreditRisk$^+$ is especially parsimonious. In exchange for the limitations of a two-state model, the model delivers the advantage of an analytic solution for the loss distribution. Calculation of percentile values of this distribution is consequently much faster than in models based on Monte Carlo simulation. Nonetheless, the calculations still are not instantaneous, even when coded in a compiled language such as C. More importantly, the standard recurrence algorithm for determining value-at-risk is fragile. For adequate recognition of sectoral and geographic concentrations, risk managers may require dozens of systematic risk factors in the model specification. At many institutions, portfolios contain ten thousand or more obligors. Yet I find that the standard algorithm often breaks down severely when modeling as few as seven systematic risk factors and five thousand obligors.

This paper addresses this problem using an alternative representation of the loss distribution. In the standard treatment of the model, the loss distribution is characterized by its probability generating function. Taylor series expansion of the pgf via the recurrence relation becomes the natural way to solve for VaR. I show that the cumulant generating function provides a more robust and computationally efficient representation of the loss distribution. At the very least, the cgf permits extremely fast and accurate computation of the moments of the loss distribution. I show how the moments provide a simple and effective diagnostic for the standard recurrence algorithm.

More interestingly, the availability of a tractable cgf implies that we can directly estimate the CreditRisk$^+$ loss distribution via saddlepoint approximation. Saddlepoint approximation offers several advantages over the standard recurrence-relation algorithm. First, it is extremely fast. Indeed, it is virtually instantaneous, even when there are many sectors and many borrowers, and even when coded in a higher-level language such as MATLAB. Second, the algorithm is more robust in practical application. It is not always accurate for small portfolios (say, 500 obligors or less), but for such portfolios the standard recurrence algorithm performs quite well, so there is no need for an alternative. For large portfolios, saddlepoint approximation is consistently accurate regardless of the complexity of the risk-factor structure.

Finally, as a loss distribution in its own right, the saddlepoint approximation may actually be better behaved in practice than the analytical loss distribution. As models of credit risk become more integrated into the management systems of financial institutions, lending decisions are increasingly driven by capital allocation. At the margin, a loan's contribution to capital is the difference between VaR for the portfolio including and excluding the loan. Because CreditRisk$^+$ is a discrete model, percentile values change in small jumps. For investment grade loans of moderate size, small changes in the choice of standardized unit of loss can produce very large relative changes in the estimated marginal tail contribution.[13] The finer the discretization of loss exposures, the more accurately can we measure marginal tail contributions, but the more fragile the standard solution algorithm. Saddlepoint approximation does not require discretization of the loss exposures, and so avoids the issue entirely. Indeed, because it implies a continuous loss distribution that changes continuously in all model parameters, all sorts of comparative statics are well-behaved.

The value of credit risk modeling lies more in the discipline it imposes on an institution's lending practices and allocation of risk capital than in the precise values of model outputs. In practice, banks are interested in measuring very high quantiles of credit loss over a long horizon (e.g., 99.8% over one year), and individual credit losses are themselves low probability events. As an empirical matter, the distributional and functional form assumptions that drive credit value-at-risk models can never be meaningfully validated.[14] Thus, a good approximation to a benchmark model may be no less valid than the model itself. It derives its discipline and legitimacy from the underlying benchmark, yet can offer significant computational benefits. In practical application, saddlepoint approximation of CreditRisk$^+$ preserves the essential characteristics of the analytical version of the original model, but usefully smooths over the jumps in the original discrete model and avoids its computational pitfalls.

A Automatic calculation of derivatives of ψ_k

The recurrence relation for differentiation of $V_{j,k}(z)$ can easily be implemented in *Mathematica*. This provides a convenient and accurate tool for checking equations (14b)–(14d) and generating higher derivatives of ψ_k. The `PsiDeriv` function is defined by:

```
(* Mathematica code courtesy of Gary Anderson. *)
Unprotect[Derivative]
Derivative[0,1][V][j_Integer,z_Symbol] := V[j+1,z] + V[j,z]*V[1,z]
Protect[Derivative]
PsiDeriv[j_Integer,z_Symbol] := tau*Simplify[D[V[1,z],{z,j-1}]]
```

To check, say, the third derivative $\psi_k'''(z)$, enter `PsiDeriv[3,z]` at the *Mathematica* input prompt. *Mathematica* returns:

```
                   3
tau ( 2 V[1, z]  + 3 V[1, z] V[2, z] + V[3, z] )
```

[13] Intuitively, small changes in the standardardized unit of loss shift the location of the "jumps" in the loss distribution, and thus can cause discrete changes in the quantiles of the distribution. This measurement issue is most serious (a) the coarser the discretization of exposure sizes, (b) the smaller the variances of the systemic risk factors, and (c) the higher the credit quality of the marginal asset.

[14] Gordy (2001) examines the fundamental role of embedded distributional and functional form assumptions.

B Bounds on z_k^*

The strict convexity of $\mathcal{Q}_k(z)$ and the definition $\mu_k = \sum_i w_{ik}\bar{p}_i$ implies that

$$\mathcal{Q}_k(z)/\mu_k = \sum_i \frac{w_{ik}\bar{p}_i}{\mu_k} \exp(\nu_i z) < \exp\left(z \sum_i \frac{w_{ik}\bar{p}_i}{\mu_k}\nu_i\right).$$

By definition, $\mathcal{Q}_k(z_k^*) = \mu_k + \tau_k = \mu_k/\delta_k$, so

$$\frac{1}{\delta_k} < \exp\left(z_k^* \sum_i w_{ik}\bar{p}_i\nu_i/\mu_k\right),$$

which implies

$$z_k^* < -\log(\delta_k) \Big/ \frac{1}{\mu_k} \sum_i w_{ik}\bar{p}_i\nu_i. \tag{20}$$

Since $\mathcal{Q}_k(z)$ is increasing and $\mathcal{Q}_k(0) = \mu_k$, we also have $0 < z_k^*$.

Now let k denote the sector for which $z^* = z_k^*$. For this sector, $\lim_{z\to z^*} V_{1,k}(z) = \infty$. As $\psi'(z) \geq \psi_k'(z) = \tau_k V_{1,k}(z)$, we have $\lim_{z\to z^*} \psi'(z) = \infty$. We also know that $\psi'(0)$ equals the mean of the loss distribution. Therefore, for any $y > \mathrm{E}[Y]$, there must exist a $\hat{z} \in (0, z^*)$ such that $\psi'(\hat{z}) = y$.

C Tail behavior of the CreditRisk$^+$ loss distribution

From the form of equation (5) for the pgf $\mathcal{G}(z)$, we see that loss can be represented as the sum of independent compound negative binomial random variables with parameters $(1-\delta_k, \tau_k)$. As shown in Panjer and Willmot (1992, §10.2), as $y \to \infty$, the tail of the cdf of a compound negative binomial variable Y can be approximated as

$$1 - F(y) \approx C(y)y^\gamma \exp(-\kappa y) \tag{21}$$

for some $\gamma \in \Re$ and $\kappa > 0$ and where $C(y)$ is slowly varying at infinity. By Embrechts, Klüppelberg and Mikosch (1997, Theorem 3.4.5), F is in the maximal domain of attraction of the Gumbel distribution if

$$\lim_{u\to\infty} \frac{1 - F(u + y/\kappa)}{1 - F(u)} = \exp(-y)$$

for all $y > 0$. From equation (21) we have for large u

$$\frac{1 - F(u + y/\kappa)}{1 - F(u)} \approx \frac{C(u + y/\kappa)}{C(u)} \frac{(u + y/\kappa)^\gamma}{u^\gamma} \exp(-\kappa y/\kappa) \approx \frac{C(u + y/\kappa)}{C(u)} \exp(-y).$$

By definition, a function that is slowly varying at infinity has $C(u + y/\kappa)/C(u) \to 1$ as $u \to \infty$. Therefore, F is MDA Gumbel, which implies that there exists a $\beta > 0$ such that G converges in the tail to an exponential distribution with parameter β (see Embrechts et al. 1997, Theorem 3.4.13). The same tail behavior applies to a sum of independent compound negative binomial variables.

References

Bürgisser, Peter, Alexandre Kurth, and Armin Wagner, "Incorporating Severity Variations into Credit Risk," *Journal of Risk*, 2001, *3* (4), 5–31.

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

Embrechts, Paul, Claudia Klüppelberg, and Thomas Mikosch, *Modelling Extremal Events* number 33. In 'Applications of Mathematics.', Berlin: Springer-Verlag, 1997.

Feuerverger, Andrey and Augustine C.M. Wong, "Computation of Value-at-Risk for Nonlinear Portfolios," *Journal of Risk*, Fall 2000, *3* (1), 37–55.

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

____ , "What Wags the Tail? Identifying the Key Assumptions in Models of Portfolio Credit Risk," in Carol Alexander, ed., *Mastering Risk*, Vol. 2: Applications, Financial Times Prentice Hall, 2001.

Jensen, Jens Ledet, *Saddlepoint Approximations*, Oxford: Clarendon Press, 1995.

Martin, Richard, Kevin Thompson, and Christopher Browne, "Taking to the Saddle," *Risk*, June 2001, pp. 91–94.

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

Press, William H., Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery, *Numerical Recipes in C*, Cambridge University Press, 1992.

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

Wood, Andrew T.A., James G. Booth, and Ronald W. Butler, "Saddlepoint Approximations to the CDF of Some Statistics with Nonnormal Limit Distributions," *Journal of the American Statistical Association*, June 1993, *88* (422), 680–686.