Enhancing CreditRisk +

Götz Giese [*]

We discuss the CreditRisk+ methodology from the perspective of the moment generating function of the credit factors. This representation lends itself to a new recursion formula for the portfolio loss distribution that is more accurate and considerably faster, particularly for large portfolios. We discuss how the model can be extended to incorporate correlations between risk factors and derive the general formula for exact VaR contributions in this modelling framework.

Since its introduction in 1997 [CSFB], CreditRisk+ (CR+) has become one of the most widely used credit portfolio models. Among its advantages is that the portfolio loss distribution can be computed analytically such that Monte Carlo simulation can be avoided. However, as has been pointed out by Gordy [3], the standard recursion relation for the loss distribution in CR+, which goes back to Panjer, [8] tends to be numerically unstable for large portfolios and many risk factors. Gordy proposes to use a saddlepoint approximation to circumvent these problems. In this article, we present a new way to calculate the true loss distribution, which is numerically stable and accurate, even under rather extreme conditions. The new method is also considerably faster than the Panjer recursion.

Our approach centres on the evaluation of the moment generating function (mgf) of the risk factor distribution. We show that the distributional assumptions in standard CR+ can be relaxed without loosing analytical tractability. We introduce a particular factor distribution ("compound gamma") that allows representing a more complex default correlation structure. We discuss how the parameters of the distribution can be fitted to an externally given covariance matrix. For a test portfolio, we compare the new model with standard CR+ and an alternative approach proposed by Bürgisser et al. [1].

In the last part of the paper, we generalise a result of Haaf and Tasche [4] on the risk-adjusted breakdown of portfolio VaR to obligor level. We show that risk contributions are linked to the partial derivatives of the mgf. With the example of the compound gamma model, we demonstrate that risk contributions from approximations (variance breakdown, saddlepoint method) can significantly deviate from the exact results.

CR+ revisited

Let us first summarise the main model ingredients of the CR+ framework. CR+ is a discrete model, which means that after defining a suitable discretisation constant ΔL the adjusted exposure E_A of each obligor A is mapped to a (positive) integer ν_A . The term "adjusted exposure" refers to the amount of money that will actually be lost in case of obligor default, including collateral, netting and (deterministic) recovery rate. For each obligor we require an estimate of the average default probability p_A for the time horizon considered (typically one year). Default of obligor A is described by a random variable N_A that can take on integer values. (Note that this implies– at least theoretically – that multiple defaults can occur over the time horizon.) The total loss of a portfolio of obligors is therefore represented by the integer random variable

$$L = \sum_A \nu_A \cdot N_A \quad .$$

Our aim is to calculate the portfolio loss distribution in terms of its probability generation function (pgf), which is the power series

$$G(z) = \sum_{n=0}^{\infty} P(L = n) \cdot z^n \ .$$

Here $|z| < R$ is a formal variable within the radius of convergence, $R \geq 1$, and $P(L = n)$ is the probability of L being equal to n. In a next step, we introduce a set of K systematic risk factors γ_k , which represent potential future states of the economy. For the moment, we do not further specify the factor distribution. We only require the variables to be non-negative and normalised,

(1) $$E^{\gamma}(\gamma_k) = 1 \ , \ \ \gamma_k \geq 0 \ \ \text{for} \ \ k = 1,\ldots,K$$

[*] Commerzbank AG, Kaiserplatz 1, D-60261 Frankfurt (e-mail: goetz.giese@commerzbank.com)

with covariance matrix

(2) $\mathrm{Cov}(\vec{\gamma})_{kl} = \sigma_{kl}$.

The risk factors are the drivers of correlated changes in the individual default rates. In the CR+ framework, one assumes the linear relationship

(3) $$p_A(\vec{\gamma}) = p_A \sum_{k=1}^{K} g_k^A \gamma_k \,, \qquad \sum_{k=1}^{K} g_k^A = 1$$

where $p_A(\vec{\gamma})$ is the default rate of obligor A *conditional* on a future state of the world as described by $\vec{\gamma}$.The standard interpretation of the systematic risk factors is to associate the γ_k with (normalised) sector default rates. In this interpretation, relative changes of individual default rates are assumed to be the same in the same industry sector. In a more general setting, the γ_k may represent forecasts of macroeconomic factors and the factor weights g_k^A will be estimated from a stochastic default model.

Conditional on $\vec{\gamma}$ the default variables N_A are assumed as independent and Poisson distributed with intensity $p_A(\vec{\gamma})$. Under these assumptions the *conditional* pgf of the portfolio loss can be easily calculated as

(4) $$G_\gamma(z) = \exp\left(\vec{\gamma} \cdot \vec{P}(z)\right)$$

where the dot denotes the inner product of vectors. $\vec{P}(z)$ is the K-dimensional vector of "sector polynomials"

(5) $$P_k(z) = \sum_A g_k^A \, p_A \left(z^{\nu_A} - 1\right),$$

which provides a condensed representation of the portfolio structure. Default correlation between obligors only arises through the dependence on the common set of risk factors. Using the linear relationship for the Poisson intensities, eq. (3), and averaging over $\vec{\gamma}$, we calculate the default correlation between two obligors A and B that are 100% assigned to sectors k_A and k_B as[1]

(6) $$\rho_{AB}^{Def} = \sqrt{\frac{p_A p_B}{\left(1 + p_A \sigma_{k_A k_A}\right)\left(1 + p_B \sigma_{k_B k_B}\right)}} \, \sigma_{k_A k_B} \approx \sqrt{p_A p_B} \, \sigma_{k_A k_B} \, .$$

Likewise, we obtain the unconditional pgf of the loss distribution by averaging $G_\gamma(z)$ over the factor distribution:

(7) $$\begin{aligned} G(z) &= E^\gamma\left(\exp\left(\vec{\gamma} \cdot \vec{P}(z)\right)\right) \\ &= M_\gamma\left(\vec{t} = \vec{P}(z)\right) \end{aligned}$$

Here $M_\gamma(\vec{t}) = E^\gamma\left(\exp(\vec{\gamma} \cdot \vec{t})\right)$ is the moment generating function (mgf) of the multivariate factor distribution. In other words, the pgf of the loss distribution is the mgf of the risk factor distribution, evaluated at the particular point $\vec{t} = \vec{P}(z)$. We will exploit this useful relationship in the next paragraph to derive a new recursion relation for the loss distribution.

In standard CR+, risk factors are assumed as independent and gamma distributed. Since the mgf of the univariate gamma distribution with mean 1 and variance σ_{kk} is $\left(1 - \sigma_{kk} t\right)^{-\frac{1}{\sigma_{kk}}}$, the pgf of the loss distribution can be written as

(8) $$G^{CR+}(z) = \exp\left(- \sum_{k=1}^{K} \frac{1}{\sigma_{kk}} \log\left(1 - \sigma_{kk} P_k(z)\right) \right).$$

New recursion scheme

Eq. (7) lays the ground for a very convenient numerical calculation of the pgf. As long as the mgf of the factor distribution has a simple analytical form, the evaluation of the rhs boils down to performing nested operations on

[1] Eq. (6) can be easily generalised to the case that obligors are assigned to more than one sector.

power series. More specifically, we represent $\vec{P}(z)$, G(z), and all other power series that occur as intermediate results in the evaluation as polynomials of fixed degree N_{max}. To evaluate expressions like Eq. (8), we have to define operators (addition, subtraction) and functions (exponential and logarithm) on polynomials, which can be very elegantly implemented in object-oriented programming languages such as C++. The exponential function can be calculated by a simple recursion relation, which follows from the power series expansion of exp(z). Let $P(z) = \sum_{n=0}^{N_{max}} p_n z^n$ and $Q(z) = \sum_{n=0}^{N_{max}} q_n z^n$ be polynomials of degree N_{max}. For $P(z) = \exp(Q(z))$ one finds the recursion relation (cf. [2])

$$p_0 = \exp(q_0),$$
$$p_n = \sum_{j=1}^{n} \frac{j}{n} q_j p_{n-j} \quad n > 0.$$

Solving for q_n yields the recursion relation for the logarithm of a polynomial (assuming $p_0 > 0$)

$$q_0 = \log(p_0),$$
$$q_n = \frac{1}{p_0}\left(p_n - \sum_{j=1}^{n-1} \frac{j}{n} q_j p_{n-j}\right) \quad n > 0.$$

Since these recursion schemes do not break off at N_{max}, the choice of this parameter determines the truncation error of the calculation[2]. The number of computation steps for the nested evaluation of eq. (8) is of order $O(K\cdot N^2_{max})$, which can be further reduced due to the fact that the maximum degree of the sector polynomials is ν_{max}, the maximal exposure size.
This algorithm has to be contrasted with the standard solution technique, known as Panjer recursion[8], which consists of representing $\frac{d}{dz}\log\left(G^{CR+}(z)\right)$ as the rational polynomial

$$\text{(9)} \qquad \frac{A(z)}{B(z)} = \sum_{k=1}^{K} \frac{\frac{d}{dz}P_k(z)}{1-\sigma_{kk}P_k(z)}.$$

Solving for A(z) and B(z) involves (K^2-1) multiplications of polynomials so that the overall computation time is of order $O(K^2N^2_{max})$, which is by a factor of K slower than the nested calculation approach. ($K \geq 50$ is quite common in practical applications.) More importantly, it is well known [3] that the repeated multiplication of polynomials in eq. (9) can lead to significant round-off errors, particularly if K, N_{max}, and the number of obligors are large. The nested calculation algorithm, on the other hand, is insensitive to this problem, since much less numerically "dangerous" floating point operations are necessary. We routinely run the algorithm on portfolios consisting of more than 10^6 obligors with K = 65 factors and N_{max} = 10,000 and have never observed any numerical instabilities. Figure 1 provides a practical example and illustrates that sometimes the Panjer may not even be able to correctly determine the 95% loss quantile.
Another advantage of the nested calculation approach is that it can be just as easily applied to extensions of the standard model, as will be discussed in the next section.

Integrating factor correlations

We now want to include correlations between risk factors into the model. We therefore consider a more general factor distribution with individual variances σ_{kk} as before and uniform covariance $\hat{\sigma}^2$. This modelling leaves the level of default correlation within sectors unchanged, but introduces additional correlation between obligors in different sectors (cf. eq. (6)).
The proposed distribution has a simple, analytically tractable mgf, such that we can still use eq. (7) to calculate the loss distribution. It is a mixture distribution, where *conditional* on a positive random variable S, each γ_k is independently gamma distributed with shape parameter

[2] Obviously, N_{max} should correspond to a portfolio loss far out in the tail of the distribution. An upper bound can be obtained from Chebyshev's inequality (cf. Merino et al. [6]). We monitor the truncation error by comparing the first moments of the loss distribution with their theoretical values, which can be derived analytically from the input parameters. In practical applications, $N_{max} = 5\cdot\nu_{max}$ (ν_{max} being the maximal exposure size) is usually a good starting point for diversified portfolios.

$$\hat{\alpha}_k(S) = S\,\alpha_k, \quad \alpha_k > 0$$

and constant scale parameter β_k. The common scaling factor S itself is also modelled as gamma distributed with $\alpha = \frac{1}{\hat{\sigma}^2}$ and $\beta = \hat{\sigma}^2$ such that

$$E^S(S) = 1 \text{ and } \mathrm{Var}(S) = \hat{\sigma}^2.$$

Here $E^S(\ldots)$ denotes the expectation taken over the probability measure of S. Averaging the conditional distribution of the γ_k over S, we obtain a K-dimensional distribution, which we call "compound gamma distribution". The joint probability density function (pdf), $f_\gamma^{CG}(\vec{\gamma})$, is therefore given by

$$f_\gamma^{CG}(\vec{\gamma}) = \int_0^\infty ds\, g_{\hat{\sigma}^{-2},\hat{\sigma}^2}(s) \prod_{k=1}^{K} g_{\hat{\alpha}_k(s),\beta_k}(\gamma_k)$$

where $g_{\alpha,\beta}(x)$ denotes the pdf of the univariate gamma distribution with shape parameter α and scale parameter β. Although it seems difficult to further simplify this expression, calculating the mgf and the cumulants of the distribution is straightforward. For the first and second cumulant one obtains

$$E(\gamma_k) = E^S\left(E^\gamma(\gamma_k \mid S)\right) = E^S(\hat{\alpha}_k(S)\,\beta_k) = \alpha_k \beta_k$$

and similarly

(10) $$\sigma_{kl} = \delta_{kl}\beta_k + \hat{\sigma}^2 \qquad \beta_k \geq 0,$$

which means that the β_k contribute only to the respective diagonal elements. (δ_{kl} denotes Kronecker's delta.) From the normalisation condition (eq. (1)) we get

(11) $$\alpha_k = \frac{1}{\beta_k}$$

such that the entire distribution is determined by $\hat{\sigma}$ and the K parameters β_k. The mgf of the compound gamma distribution reads

(12) $$M_\gamma^{CG}(\vec{t}) = \exp\left\{-\frac{1}{\hat{\sigma}^2}\log\left(1 + \hat{\sigma}^2 \sum_{k=1}^{K} \frac{1}{\beta_k}\log(1 - \beta_k t_k)\right)\right\}.$$

This representation of the mgf can be easily plugged into eq. (7) to calculate the portfolio loss distribution. The computation time is still of order $O(KN_{max}^2)$.

An interesting way to look at the compound gamma distribution is to consider it as an "interpolation" between the independent gamma distribution ($\hat{\sigma} = 0$) and a single-factor gamma distribution (all $\beta_k = 0$). In the latter case the multivariate distribution simply consists of K replicas of a single gamma distributed variable with variance $\hat{\sigma}^2$. For intermediate parameter values, the marginal distribution has to be computed numerically by Fourier inversion of the characteristic function. For the k^{th} factor, the characteristic function reads

(13) $$C_k^{CG}(t) = \left[1 + \frac{\hat{\sigma}^2}{\beta_k}\log(1 - i\beta_k t)\right]^{-\frac{1}{\hat{\sigma}^2}}.$$

As can be seen from Figure 2, the differences between gamma and compound gamma distribution are rather small in the parameter range of practical interest ($\sigma_{kk} < 1$). In general, the compound gamma distribution is slightly more fat-tailed. Test calculations comparing standard CR+ with a model with independent compound gamma distributed factors did not show a significant effect on the portfolio loss distribution, when both distributions were calibrated to the same variance structure. This is in line with the findings of Koyluoglu and Hickman [5]. In the end, the specific choice of the marginal factor distribution is mainly a matter of mathematical convenience.

Fitting to a given covariance structure

In practical applications, one typically estimates a covariance matrix V_{kl} from time series data, e.g. from a regression analysis of historical sector default rates. We now need to fit the K+1 parameters of the compound gamma distribution to this (generally more complex) covariance structure. We determine the parameters by matching the first and second moments of the loss distribution. It suffices to focus on the loss variance, since eq. (1) already guarantees matching of the first moment. Loss variance can be split into a portfolio-specific and a systematic part. The systematic part depends on the factor distribution only through the covariance matrix V_{kl} :

$$(14) \qquad \sigma_{sys}^{Loss} = \sum_{k,l=1}^{K} EL_k V_{kl} EL_l \quad ,$$

where $EL_k = \frac{d}{dz} P_k(z=1)$ is the expected loss of sector k. Since the compound gamma distribution allows fitting the diagonal terms V_{kk} exactly, we only need to match the contribution of the off-diagonal terms in (14). Solving for $\hat{\sigma}$ leads to

$$(15) \qquad \hat{\sigma}^2 = \frac{\sum_{k \neq l} EL_k V_{kl} EL_l}{\sum_{k \neq l} EL_k EL_l} \quad .$$

This approach requires some caution because both $\hat{\sigma}^2$ and the β_k must be non-negative. Non-negativity of $\hat{\sigma}^2$ is in practice not a constraint, since the average correlation between credit risk factors is usually positive. For a large number of factors, however, and depending on the particular form of V_{kl}, $\hat{\sigma}^2$ from eq. (15) may be larger than some of the diagonal variances V_{kk}. This would require negative β_k for these factors (cf. eq. (10)). In this case, an iterative fitting procedure needs to be applied that stepwise includes diagonal terms in the numerator and denominator of eq. (15) until all parameter restrictions are fulfilled.

Model Comparison

We now compare the model with the correlation approach proposed by Bürgisser et al. [1], which consists of calculating a single-factor CR+ model with matching loss variance. (As pointed out above, the single-factor model is in fact a special case of the compound gamma model.) To highlight the differences, we consider a particular test portfolio, which is made up of K = 12 sectors each containing 3,000 obligors. Obligors in sectors 1 to 10 belong to equal parts to one of three homogenous classes with adjusted exposures $E_I = 1$, $E_{II} = 2.5$, $E_{III} = 5$ monetary units and respective default probabilities $p_I = 5.5\%$, $p_{II} = 0.8\%$, and $p_{III} = 0.2\%$. For the three obligor classes in sectors 11 and 12, we assume the same default rates, but twice as large exposures ($E_I = 2$, $E_{II} = 5$, $E_{III} = 10$). We consider a uniform factor variance $V_{kk} = 0.04$ for all factors except for sector 12, where we set $V_{1212} = 0.49$. For simplicity, we assume a uniform sector correlation $\rho = 0.1$, such that $V_{kl} = 0.1 \cdot \sqrt{V_{kk} V_{ll}}$ for the off-diagonal elements of the covariance matrix. In this set-up, sectors 11 and 12 carry the dominating part of the portfolio loss variance (9% and 58%, respectively), whereas the other sectors contribute only 3.3% each.

Table 1 compares the statistics of the loss distribution from the compound gamma model with those from standard CR+ (no factor correlation) and the Bürgisser model. As expected, the introduction of sector correlations increases the standard deviation of both the Bürgisser and the compound gamma model. The single-factor model, however, produces implausible tail behaviour: For confidence levels larger than 99% the loss quantiles from this distribution are even smaller than those of the CR+ model. Since the single-factor model replaces the heterogeneous covariance structure by an overall variance $\sigma_{sf}^2 = 0.019$, it cannot adequately capture the concentration risk in sector 12 (large exposures and large factor variance). The compound gamma model, on the other hand, fully preserves the correct default correlation within the sectors by using the original V_{kk} and employs a much smaller average inter-sector covariance $\hat{\sigma}^2 = 0.007$. As a consequence, the resulting loss distribution is consistently more fat-tailed than the other distributions.

Exact risk contributions

For a portfolio manager, perhaps the most important feature of a portfolio model is its ability to break down portfolio risk to obligor level in a risk-adjusted way. Being additive, risk contributions per obligor can be used to individually allocate economic capital. They can be aggregated across arbitrary dimensions in order to “slice and dice” through portfolio risk. If the risk measure under consideration is value-at-risk, we seek for an additive breakdown of the q% portfolio loss quantile, l_q, as calculated from the loss distribution eq. (7). We define the quantile contribution[3] QC_A of obligor A as the expected individual loss *conditional* on the portfolio loss being equal to l_q, i.e.

$$(16)\qquad \begin{aligned} QC_A &= \nu_A E\left(N_A \middle| L = l_q\right) \\ &= \nu_A \frac{E\left(N_A \cdot \mathbf{1}_{\{L=l_q\}}\right)}{P\left(L = l_q\right)} . \end{aligned}$$

Here $\mathbf{1}_B$ is the indicator function, which is equal to 1 on the event B and zero otherwise. As has been shown in a recent paper by Haaf and Tasche [4], there exists an exact analytical formula for quantile contributions in standard CR+. In the following, we generalise these results to the case of an arbitrary factor distribution respecting eq. (1).

Since the denominator in eq. (16) is directly given by the lq-th coefficient of the pgf in eq. (7), the difficult task is to evaluate the expectation in the numerator. Using the fact that conditional on γ the default variables N_A are independent Poisson variables, Haaf et al. show that (in our notation)

$$(17)\qquad E\left(N_A \cdot \mathbf{1}_{\{L=l_q\}}\right) = E^{\gamma}\left(p_A(\vec{\gamma}) \cdot P\left(L = l_q - \nu_A \,\middle|\, \vec{\gamma}\right)\right).$$

The conditional probability on the rhs is given by the $(l_q\text{-}\nu_A)$-th coefficient of the conditional pgf, $G_\gamma(z)$ (eq.(4)). For ease of notation, we introduce the "coefficient operator" $\boldsymbol{D}^{(k)}$

$$\boldsymbol{D}^{(k)} Q(z) = \frac{1}{k!} \frac{d^k}{d\, z^k} Q(z = 0)$$

which returns the k-th coefficient of an (absolutely converging) power series Q(z). Interchanging integration and differentiation, we can now transform eq. (17) to

$$(18)\qquad \begin{aligned} E\left(N_A \cdot \mathbf{1}_{\{L=l_q\}}\right) &= p_A \sum_{k=1}^{K} g_A^k \cdot \boldsymbol{D}^{(l_q-\nu_A)} E^{\gamma}\left(\gamma_k \cdot \exp\left(\vec{\gamma} \cdot \vec{P}(z)\right)\right) \\ &= p_A \sum_{k=1}^{K} g_A^k \cdot \boldsymbol{D}^{(l_q-\nu_A)} \frac{\partial}{\partial\, t_k} M_\gamma\left(\vec{t} = \vec{P}(z)\right) \end{aligned}$$

which reveals that risk contributions are related to the partial derivatives of the mgf (provided they exist). It can be shown that $G_k(z) := \frac{\partial}{\partial\, t_k} M_\gamma\left(\vec{t} = \vec{P}(z)\right)$ is also a pgf. Computing risk contributions for the portfolio therefore requires calculating K additional power series. The final formula for the quantile contribution per obligor reads

$$(19)\qquad QC_A = p_A \nu_A \frac{\sum_{k=1}^{K} g_A^k \cdot \boldsymbol{D}^{(l_q-\nu_A)} G_k(z)}{\boldsymbol{D}^{(l_q)} G(z)} .$$

Note that the individual default rate appears only in form of the obligor's expected loss, $p_A \nu_A$. The coefficients of the power series $G_k(z)$ and G(z) determine sector-specific "penalty factors" for large exposures, which are independent of the default rate[4].

It is interesting to compare the exact result with the formula for risk contributions in the saddlepoint approximation, as proposed by Martin et al. [6]. In our notation, this reads

$$(20)\qquad QC_A^{SP} = p_A \nu_A \cdot e^{\hat{t}\nu_A} \frac{\sum_{k=1}^{K} g_A^k \cdot G_k(z = e^{\hat{t}})}{G(z = e^{\hat{t}})} ,$$

---

[3] VaR is usually defined as *unexpected* loss on a q% confidence level. The VaR contribution of obligor A is therefore his quantile contribution minus expected loss, $VaRC_A = QC_A - p_A \nu_A$.

[4] Note that in extreme cases (large exposure and default rate) $QC_A > \nu_A$ is possible, which is a consequence of using Poisson variables as default indicators.

where $\hat{t} < \log(R)$ is the unique solution of the saddlepoint equation $\frac{d}{dt} G(e^t) = l_q \cdot G(e^t)$. For a diversified portfolio, the saddlepoint method accurately determines the tail probability $P(L \geq l_q)$ [3]. Applied to risk contributions, however, the method produces only a uniform exponential penalty factor for all obligors (all terms on the fraction in eq. (20) are constants). Comparing with the exact result we see that *all* distributions generated by the $G^{(k)}(z)$ must have the same decay behaviour should the saddlepoint result be accurate.
For the compound gamma distribution, the $G_k(z)$ have a simple form and can be easily computed by nested calculation as described above:

(21) $$G_k^{CG}(z) = \exp\left\{ \left(1 + \hat{\sigma}^2\right) \log G^{CG}(z) - \log\left(1 - \beta_k P_k(z)\right) \right\}.$$

For the test portfolio of the preceding section, Table 2 shows risk contributions aggregated by sectors using the compound gamma model. We compare exact VaR contributions with saddlepoint VaR contributions and an additive breakdown of loss variance (as proposed in [2]). Not surprisingly, the impact of large exposures (sectors 11 and 12) on tail risk is significantly underestimated by a variance-based contribution method. On the other hand, the saddlepoint approximation overestimates the contribution of the dominating sector (sector 12), which leads to a relatively high error for the small sectors 1-10. Here the saddlepoint result is more than 13% below the true value. For the tail probability, however, the saddlepoint approximation is nearly perfect (less than $3 \cdot 10^{-6}$ relative error). As pointed out before, the method is tailor-made for fitting the loss distribution, but can quite well distort the probability measure on a sub-portfolio level.

Conclusions

Generalising the standard CR+ model, we have presented a framework to determine the loss distribution and the risk contributions of a credit portfolio without approximations and numerical instabilities. The cornerstone of our approach is the mgf of the factor distribution, for which a simple analytic representation needs to be available. We have discussed a particular model variant with uniform default correlation between sectors. It would be interesting to apply our concept to other factor distributions with yet more complex dependence structure.

References

[1] P. Bürgisser, A. Kurth, A. Wagner, and M. Wolf, *Integrating correlations*, RISK July 1999, p. 57

[2] Credit Suisse Financial Products, *CreditRisk+, a credit risk management framework*, 1997, http://www.csfb.com/ creditrisk

[3] M.B. Gordy, *Saddlepoint Approximation of CreditRisk+* , Journal of Banking and Finance, Volume 26, Number 7, August 2002, pp. 1337-1355

[4] H. Haaf and D. Tasche, *Credit portfolio measurements*, GARP Risk Review July/August 2002, p.43

[5] H. Koyluoglu and A. Hickman, *Reconcilable differences*, RISK October 1998, p. 56

[6] R. Martin, K. Thompson, and C. Browne, *VaR: who contributes and how much?*, RISK August 2001, p. 99

[7] S. Merino and M. Nyfeler, *Calculating portfolio loss*, RISK August 2002, p. 82

[8] H.H. Panjer and G.E. Willmot, *Insurance risk models*, Society of actuaries, Schaumberg, IL, 1992

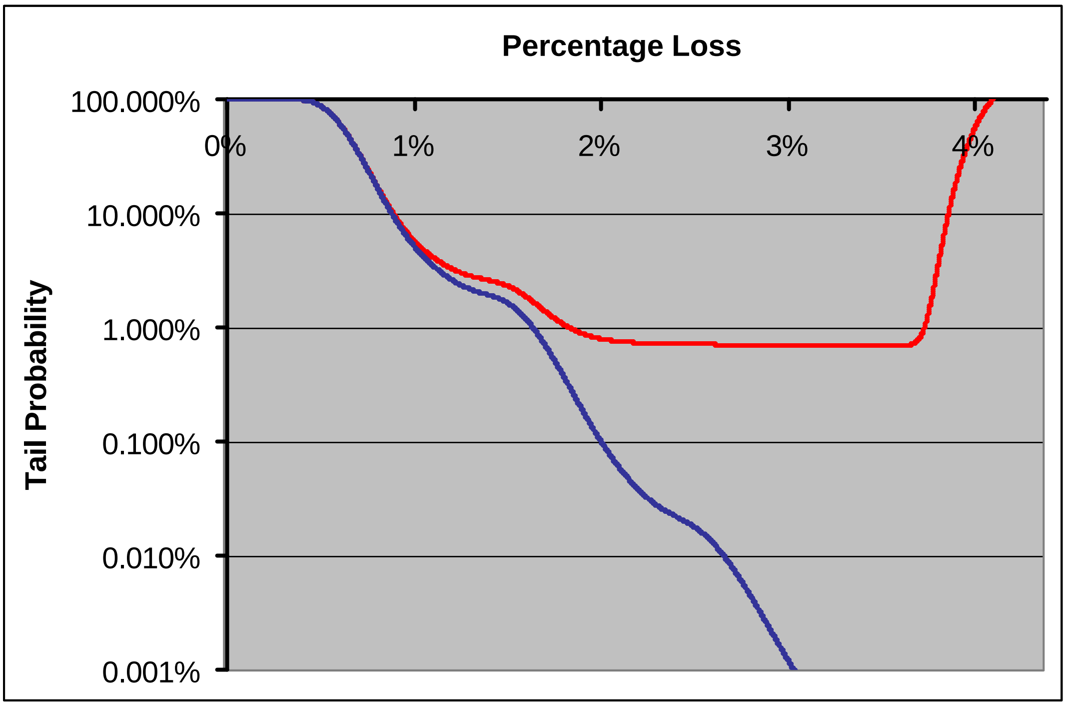

Figure 1: Tail probabilities for a very large portfolio (1.4 million obligors distributed over 65 sectors) in the standard CR+ framework. The graph demonstrates the breakdown of the Panjer recursion (red line) due to numerical instability. Numerical errors cumulate long before the 1% tail has been reached producing even negative probabilities in the far tail (area where tail probability increases). The blue distribution has been generated by nested evaluation of polynomials of degree $N_{max} = 20{,}000$.

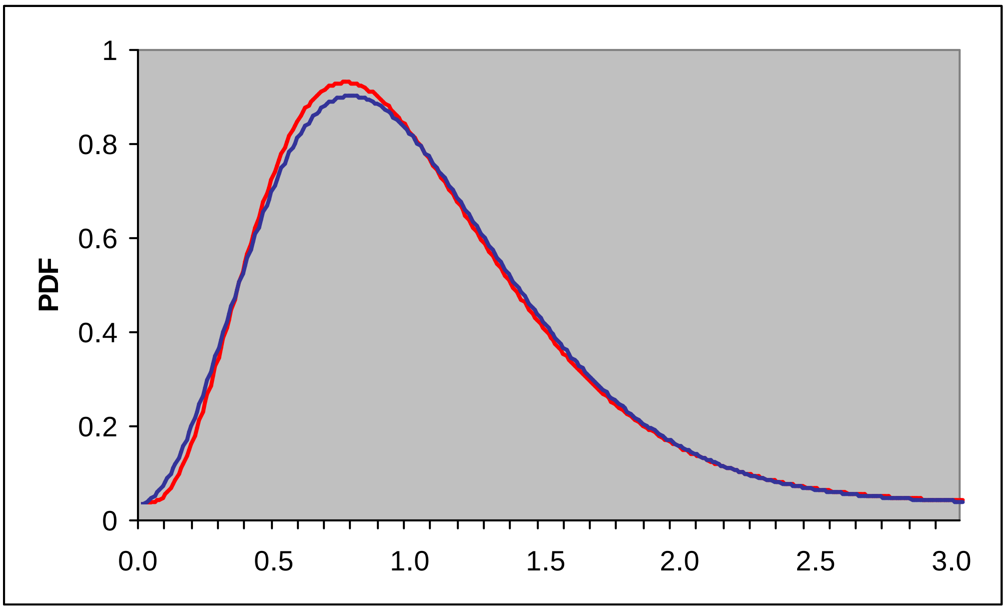

Figure 2: Comparison of the probability density functions of the gamma (red) and the compound gamma distribution (blue) with mean 1. Both distributions have the same standard deviation $\sigma = 0.5$. The parameters of the compound gamma distribution are $\hat{\sigma} = 0.3$ and $\beta = 0.16$.

	Standard CR+	Equivalent 1-Factor Model	Compound Gamma
Inter-sector Covariance	-	0.019	0.007
Expected Loss	1.00%	1.00%	1.00%
Std Deviation	0.12%	0.14%	0.14%
99% Quantile	1.36%	1.37%	1.40%
99.5% Quantile	1.42%	1.41%	1.46%
99.9% Quantile	1.55%	1.50%	1.60%

Table 1: Comparison of loss distributions from standard CR+, the single-factor, and the compound gamma model for the test portfolio described in the text. All loss statistics are quoted as percentages of the total adjusted exposure.

Sector	Variance	Saddle-point	Exact
1 ,..., 10	3.3%	1.7%	2.0%
11	9.3%	4.7%	5.4%
12	58.1%	78.2%	74.9%

Table 2: Aggregated risk contributions (in percent) for the test portfolio. Contributions to loss variance are being compared with the saddlepoint approximation and the exact formula for the risk-adjusted breakdown of VaR (on a 99.9% confidence level).