Analytical techniques for synthetic CDOs and credit default risk measures

A. Antonov*, S. Mechkov†, and T. Misirpashaev‡

NumeriX

May 23, 2005

Abstract

Pricing and risk management of synthetic CDOs and risk management of credit portfolios are closely related problems as both require modeling of the same distribution of portfolio loss. The valuation of a single tranche CDO is equivalent in complexity to the calculation of credit default VaR for a portfolio of single name entities, while the valuation of CDO^2 (CDO-squared) is a task closely related to the calculation of credit default VaR for a portfolio of single tranche CDOs. We examine the analytical techniques developed for credit portfolio problems with a view to CDO applications and find that the saddlepoint method works better than the alternatives, leading to a new, fast technique for CDO^2 pricing and hedging.

1 Introduction

The total loss due to defaults $L(T)$ which will be accumulated as a result of holding a portfolio of risky assets until the time horizon T is a major unknown. The classic approach to dealing with the default risk calls for its measurement, budgeting, and attribution. The advent of credit derivatives opened an entirely new set of possibilities revolving around hedging and synthetic risk re-engineering. The search for winning financial strategies in both alleys depends upon the ability to find an adequate description of $L(T)$ in terms of a probability distribution. Depending on the particular portfolio, risk measure, or derivative instrument, the task may be as simple as estimating statistics for the distribution of portfolio loss $L(T)$ at a single time horizon or significantly more complex, involving joint distribution of loss in several correlated portfolios $\{L_p(T_i)\}$ at different time horizons $\{T_i\}$.

The typical problem of portfolio risk management is the evaluation of the value-at-risk and expected shortfall along with a meaningful decomposition into asset contribution. The typical risk re-engineering problem is the valuation of synthetic CDO and CDO^2 (CDO-squared)

*NumeriX Software Ltd. 28 Austin Friars 2nd Floor London EC2N 2QQ UK; antonov@numerix.com
†NumeriX Software Ltd. 28 Austin Friars 2nd Floor London EC2N 2QQ UK; meshkov@numerix.com
‡NumeriX LLC 60 E. 42nd St., Suite 2434, New York, NY 10165 USA; missir@numerix.com

instruments along with hedge ratios to single name credit spreads. There are notable distinctions in the interpretation of the probability distributions used in VaR-type problems as opposed to CDO-type problems. VaR and other portfolio risk measures should be computed in the so-called real world measure while CDOs and other credit derivatives should be valued in the risk neutral measure; see e.g. Bluhm et al. (2003). A simplified but useful way to state the distinction between the two measures is in terms of the preferred source of default probabilities data. For the real world measure the adequate source would be historical default frequencies. For the risk neutral measure the preferred source is tradable single name credit sensitive instruments, such as credit default swaps and bonds. Default probabilities implied by tradable instruments appear to be systematically higher than those computed from historical data. The deeper reasons for this discrepancy is a subject of intense ongoing research; see, e.g., recent works by Hull et al. (2005) and Berndt et al. (2005).

A necessary ingredient in every model of portfolio loss is the model of default correlation. There is a belief that here too a distinction can be drawn between real world and risk neutral correlations, however it is very difficult to quantify this belief, especially because the industry has yet to settle on a satisfactory correlation model. Furthermore the very presence of correlations in a model essentially blurs the line between real world and risk neutral valuations (a point strongly emphasized by Rebonato (2004)), which is waiting to be properly resolved in a future, better theory. Lacking such a theory we will adhere to the mainstream copula models of correlations and focus on the practical side of the computation of risk measures and expected payoffs. In doing so we will be exploring mathematical equivalence of VaR-type and CDO-type problems without specifying whether the measure is real world or risk neutral and assuming that the default probabilities and correlations are always chosen in accordance with the specific application of the results.

The mathematical equivalence is obvious from the fact that the stop-loss option $E[(L-K)^+]$, which is the key building block of synthetic CDO valuation, is also the key term in the expression for the expected shortfall, $E[L|L \geq K] = K + E[(L-K)^+]/P[L \geq K]$. The denominator in this expression for the shortfall is the tail probability $P[L \geq K]$, which is the inverse function to the value-at-risk and is equal to the negative of the derivative of the stop-loss option with respect to strike

$$P[L \geq K] = -\frac{\partial E[(L-K)^+]}{\partial K}. \tag{1}$$

Introducing the cumulant generating function $\mathcal{K}(\xi) = \ln E[\exp(\xi L)]$ (further referred to as CGF) and evaluating the stop-loss option and the tail probability as an inverse Laplace transform, we get expressions that differ only by the power of ξ in the denominator,

$$E[(L-K)^+] = \frac{1}{2\pi i}\int_{c-i\infty}^{c+i\infty} \frac{\exp(\mathcal{K}(\xi)-\xi K)}{\xi^2} d\xi, \tag{2}$$

$$P[L \geq K] = \frac{1}{2\pi i}\int_{c-i\infty}^{c+i\infty} \frac{\exp(\mathcal{K}(\xi)-\xi K)}{\xi} d\xi \tag{3}$$

(here $c > 0$ is any positive number and the contour of integration is parallel to the imaginary axis). This suggests that any analytical method capable of computing the stop-loss should be adaptable to the tail probability and vice versa, even though in most cases it would be impractical to compute one from the other by means of numerical integration or differentiation.

In this paper, we review the methods that have been used for VaR-type problems and CDO-type problems and explore unused opportunities, mostly in the direction of adaptation of the methods previously used predominantly for VaR-type problems to CDOs. Our specific attention is on the normal proxy approach and the saddlepoint method, the latter being treated as an advanced refinement of the former. The results for a single tranche CDO are both analytically appealing and computationally efficient. For CDO^2 the result can be written down in a closed form but calls for further simplifications to ease numerical evaluation. We discuss the conditional decoupling of multivariate normal proxy as a means of speeding up the calculation.

The rest of the paper is organized as follows. We start the thread of analytics in Sect. 2 by a discussion of copula functions based correlation models, then take a pause in Sect. 3 to state precise definition of risk measures and CDO cashflows. This section also includes a simple reasoning as to why a single implied correlation number is bound to exhibit a pathological behavior for mezzanine CDO tranches in almost any reasonable correlation model and how base correlations manage to escape the argument and offer at least a partial solution. The analytical thread is resumed in Sect. 4 where we tackle the simpler case of single tranche CDOs and risk measures for portfolios without CDOs. The quadratic decomposition of the CGF around 0 gives rise to the normal proxy model. The saddlepoint is introduced as a better point around which to expand the CGF. Building on the results pioneered by Martin, Thompson, and Browne in the series of publications (2001a,b,c), we derive the corrections to the saddlepoint approximation for the tail probability, expected shortfall, and stop-loss, and give explicit expressions for the decomposition of expected shortfall in the leading saddlepoint approximation. This is followed by a discussion of Gordy's granularity adjustment (2003) to show how it fits in the framework of CGF expansions. The numerical comparison of the granularity adjustment, normal proxy integration, and saddlepoint method shows that the saddlepoint emerges as a clear winner. The section is concluded by a brief digression into Monte Carlo simulation with importance sampling, which offers an intuitive understanding of the reason for saddlepoint effectiveness. Sect. 5 is devoted to applications of the same ideas to CDO^2. A dramatic reduction in the dimensionality of integration is possible however further simplifications are needed to get rid of the simulation completely. We pursue one possible approach based on the factor decorrelation of portfolio loss and give numerical results before concluding in Sect. 6.

2 Models of correlated defaults

Let N be the number of the assets in the underlying portfolio, each of which is in good standing as of the valuation date $t = 0$ but has a non-zero probability of default by the horizon date $t = T$. The standard approach to joint default modeling consists of two steps. In the first step independent single name default models are defined. In the second step a joint distribution of defaults is formed to extend the marginal distributions of defaults fixed in step one. The single name credit default model for asset a consists of its survival probability curve $s_a(t)$ and a model for loss-given-default. The relevant output that summarizes the loss-given-default model for the purpose of credit portfolio modeling is the weight w_a of each asset in the decomposition of the portfolio loss in terms of asset default indicators $U_a(T)$,

$$L(T) = \sum_a w_a U_a(T). \tag{4}$$

In the case of several portfolios, a separate weight $w_{a,p}$ needs to be introduced for loss exposure of the asset a in each portfolio p,

$$L_p(T) = \sum_a w_{a,p} U_a(T). \tag{5}$$

The valuation of synthetic CDO tranches that involve senior slices of the portfolio is also affected by expected values of tranche amortization due to recovery. This requires a second set of weights, w'_a, to express the random recovery as

$$R(T) = \sum_a w'_a U_a(T) \tag{6}$$

(and similarly in the case of several portfolios). In the popular fractional recovery model of loss-given-default, the parameter is the recovery rate r_a, which relates the loss and recovery rates to the asset notional A_a as $w_a = (1-r_a)A_a$, $w'_a = r_a A_a$ but generally the loss and recovery weights need not be constrained in this way. Because the problem of recovery distribution is an exact duplicate of the problem of loss distribution we do not need to develop any specific techniques for recovery and will focus on loss.

The current industry standard of introducing the correlations is the Gaussian copula model, Li (2000). An operational definition of this model can be given by describing how to simulate the joint distribution of default status of the N assets. We introduce an $N \times N$ matrix of copula model correlations, ρ_{ab}. On each simulation path a random vector $(Z_1, \ldots, Z_N)$ is drawn from the multivariate normal N-dimensional distribution with mean 0 and covariance matrix ρ_{ab} so that the density of this vector is given by

$$p(Z_1, \ldots, Z_N) = \frac{1}{\sqrt{(2\pi)^N \det \rho}} \exp(-\frac{1}{2}\mathbf{Z}^T \rho^{-1} \mathbf{Z}). \tag{7}$$

For any time horizon the generated random numbers are compared with the threshold barriers, $b_a(T) = \mathcal{N}^{-1}(1 - s_a(T))$. If $Z_a \leq b_a$ we declare that the asset a has defaulted by the time T, otherwise that it has remained alive. Here $\mathcal{N}(x) = (1/2\pi)\int_0^x \exp(-0.5t^2)dt$ is the $N(0,1)$ cumulative distribution function, and $\mathcal{N}^{-1}$ is the inverse function to $\mathcal{N}$. The explicit expression for the default indicator variable in this model is $U_a(T) = \theta(b_a(T) - Z_a)$, where $\theta(x)$ is the step function, equal to 1 for $x \geq 0$ and 0 otherwise.

The full correlation matrix is unwieldy, obscures the source of correlations as coming from specific common factors, and also makes non-Gaussian extensions less transparent. A factor representation is often superior, in which the correlation matrix is decomposed as $\rho_{ab} = \sum_i \beta_a^{(i)} \beta_b{}^{(i)}$. The numbers $\beta_a^{(i)}$ are called factor loadings. The factorized version of the fundamental simulation for the copula model takes the form

$$Z_a = \sum_i \beta_a^{(i)} X_i + \sqrt{1 - \beta_a^2} Y_a, \tag{8}$$

with each component X_i of the global factor and drivers Y_a of individual risk factor being independent $N(0,1)$-distributed random variables. The factor formulation of the copula model suggests a two-stage procedure for the calculation of expectations. At the first stage the average is taken over the individual risk factors $\{Y_a\}$ conditionally on a realization of the global factor $\mathbf{X}$. This is facilitated by independence of the individual risk factors. The second stage is the Gaussian integration over the global factor (which will be denoted $E_{\mathbf{X}}[...]$). Stochastic asset weights and correlations driven by the global factor can be easily accommodated at this step as shown by Andersen and Sidenius (2005). The computational difficulty of the integration increases rapidly with the dimensionality, which is why software implementations are often restricted to the scalar global factor.

It is possible to obtain families of non-Gaussian modifications of the copula model using non-normal distributions for the factors in (8); see, e.g., Schönbucher (2002) and Burtschell et al. (2005). A popular case involves Student's t-distribution for the sake of its fatter tails (polynomial rather than exponential). However no dramatic improvements in the market matching ability of the model were discovered. When only one factor is used with equal elements corresponding to all assets, we get the simplest correlations model with the correlation matrix where all off-diagonal elements are equal. A simple consideration (see Sect. 3.3 below) shows that any model with a single correlation number suffers from anomalies in CDO tranche pricing but this does not prevent practitioners from using single correlation numbers as a quotational convention.

Specific applications use shortcuts to obtain the desired averages over the loss distribution directly without computing the entire probability distribution function. We will continue the discussion in terms of the probability distribution function in this section and will specify the applications in this next section.

Following the two-stage procedure we introduce conditional default probabilities $\mu_a(\mathbf{X})$ which are averages of the default indicators U_a conditional on a certain value of the global factor $\mathbf{X}$,

$$\mu_a(\mathbf{X}) = E[U_a|\mathbf{X}] = \mathcal{N}\left(\frac{b_a - \beta_a \mathbf{X}}{\sqrt{1-\beta_a^2}}\right) \tag{9}$$

The loss L conditional on the global factor $\mathbf{X}$ is a sum of independent Bernoulli random variables with two-state probability distributions $P_a(L|\mathbf{X})$,

$$P_a(L|\mathbf{X}) = \begin{cases} 1-\mu_a(\mathbf{X}) & \text{for} \quad L = 0, \\ \mu_a(\mathbf{X}) & \text{for} \quad L = w_a. \end{cases} \tag{10}$$

The probability distribution of any sum of independent random variables is the convolution of the individual distribution functions. (For any two functions $f_1(x)$ and $f_2(x)$ their convolution, denoted $f_1 * f_2$, is the function $(f_1 * f_2)(x) = \int f_1(y) f_2(x-y) dy$. The convolution of more than two functions, $f_1 * f_2 * \ldots f_N$ can be reduced to pairwise convolutions as $f_1 * (f_2 * (f_3 * \ldots * (f_{N-1} * f_N) \ldots))$ or in any other order.)

There are two ways to compute the convolution. The direct way is to use the definition and take the integrals one after another. This approach is especially convenient when all loss weights w_a are equal, or at least commensurate. In this case the distribution of loss is discrete with the support which is a grid with not too many nodes. Then the nested convolution can be implemented as a highly efficient recursive scheme, as described by Andersen et al. (2003), Hull and White (2004), and Brasch (2004). This approach becomes less efficient when the loss weights are incommensurate and the grid of possible values of loss is dense. It is particularly unsuitable for the calculation of the sensitivities to loss weights (necessary for VaR decomposition) because a small change in one of the weights gives rise to a radical change in the support of the distribution of loss.

Fortunately there is a complementary approach to computing the convolution which works better right where the direct approach fails. This approach takes advantage of either Fourier or Laplace transforms to replace the convolution by the product of Fourier or Laplace images and then return to the original space by means of the inverse transform. (Remarkably, even though this approach is analytically more involved it was introduced for synthetic CDO valuation first, and only later was it noticed that the direct convolution can be superior.) This approach also serves as the basis for various semianalytical approximations, including the central limit theorems and the saddlepoint. Here we use a version of the Laplace transform with a non-standard choice of sign in the exponent, which is however consistent with the standard definition of the conditional CGF, $\mathcal{K}(\xi|\mathbf{X}) = \ln E[\exp(\xi L)|\mathbf{X}]$, where ξ is a variable in the complex plane. The conditional CGF based on the distributions (10) computes easily,

$$\mathcal{K}(\xi|\mathbf{X}) = \sum_a \ln(1 - \mu_a(\mathbf{X}) + \mu_a(\mathbf{X}) \exp(\xi w_a)). \tag{11}$$

The conditional probability density of the loss is restored from the conditional CGF by the inverse Laplace transform

$$P(L|\mathbf{X}) = \frac{1}{2\pi i} \int_{c-i\infty}^{c+i\infty} \exp(\mathcal{K}(\xi|\mathbf{X}) - \xi L) d\xi, \tag{12}$$

where c is any positive real number.

The two-stage procedure of computing the distribution conditional on the global factor $\mathbf{X}$ and integrating over $\mathbf{X}$ after that extends to the case of the joint distribution of loss from several portfolios. Both the convolution approach and the Fourier/Laplace transform approach formally carry over to the multi-portfolio case as well. Unfortunately, the computational difficulty of the convolution approach grows dramatically because, instead of a linear grid of possible values of loss, the recursive scheme would have to operate on an M-dimensional lattice where M is the number of portfolios. This makes the numerical convolution practically unusable for more than two portfolios. The difficulties of the Laplace transform also increase substantially but the approach still holds promise. The conditional CGF in this case is a function of M complex variables, $\mathcal{K}(\xi_1, \ldots, \xi_M|\mathbf{X}) = \ln E[\exp(\sum_p \xi_p L_p)|\mathbf{X}]$, given explicitly by

$$\mathcal{K}(\xi_1, \ldots, \xi_M|\mathbf{X}) = \sum_a \ln \left(1 - \mu_a(\mathbf{X}) + \mu_a(\mathbf{X}) \exp\left(\sum_p \xi_p w_{a,p}\right)\right). \tag{13}$$

The joint probability density of loss of M portfolios is given by the multidimensional inverse Laplace transform

$$P(L_1, \ldots, L_M|\mathbf{X}) = \int_{\mathbf{c}-i\infty}^{\mathbf{c}+i\infty} \frac{\prod d\xi_p}{(2\pi i)^M} \exp(\mathcal{K}(\xi_1, \ldots, \xi_M|\mathbf{X}) - \sum_p \xi_p L_p), \tag{14}$$

where $\mathbf{c} = (c_1, \ldots, c_p)$ is any vector with real positive elements.

3 Instruments and portfolio risk measures

In this section, we pause the development of analytical methods to take stock of the applications and provide a description of the instruments and portfolio risk measures to which the models of correlated defaults are applied.

3.1 Risk measures for credit portfolio

Value-at-risk at a given confidence level α is the quantile of the loss distribution,

$$\mathrm{VaR}_\alpha(L) = \inf\{\ell \geq 0 | P[L \leq \ell] \geq \alpha\}. \tag{15}$$

For a continuous distribution this would mean that the VaR is the threshold value of the loss such that only the fraction $1-\alpha$ of the outcomes lead to loss equal to or larger than the threshold. For a discrete distribution the fraction of the outcomes above or at the level of VaR defined by Eq. (15) may occur to be smaller than $1-\alpha$ but this effect of distribution granularity is essentially irrelevant for large portfolios.

Expected shortfall at a given confidence level α is the expected value of the loss in excess of $\mathrm{VaR}_\alpha(L)$,

$$\mathrm{ES}_\alpha(L) = E[L | L \geq \mathrm{VaR}_\alpha(L)] \tag{16}$$

(Note that the normalization denominator $1-\alpha$ is already included in the definition of conditional expectation.)

To assess the contributions of individual assets to VaR and ES, it is necessary to be able to compute sensitivities to asset weights, $\partial \mathrm{VaR}_\alpha(L)/\partial w_a$ and $\partial \mathrm{ES}_\alpha(L)/\partial w_a$. Note that VaR and ES are homogeneous of degree 1 with respect to portfolio loss weights w_a, that is $\mathrm{VaR}_\alpha(L; v \cdot w_1, \ldots, v \cdot w_N) = v \cdot \mathrm{VaR}_\alpha(L; w_1, \ldots, w_N)$ and similarly for ES. It follows that

$$\mathrm{VaR}_\alpha(L) = \sum_a w_a \frac{\partial \mathrm{VaR}_\alpha(L)}{\partial w_a} \tag{17}$$

with a similar decomposition of ES into marginal contributions from the individual underlying names.

3.2 Single tranche CDO

A single tranche CDO is a synthetic basket credit instrument which involves two parties and references a portfolio of credit names. One party is the buyer of the protection, the other is the seller of the protection. Synthetic means that the cashflows are determined only by the information about the default status of obligations issued by the underlying credit names so that neither of the parties has to actually possess any of those obligations. A single tranche CDO contract defines two bounds, $k < K$, called attachment points and usually quoted as percentage points of the total original reference notional A of the underlying portfolio. The lowest tranche, $0-3\%$ or similar, is customarily called equity tranche. The highest tranche, or $30\%-100\%$, is called senior (or supersenior) tranche. The other tranches are called mezzanine tranches. The difference of the bounds $K-k$ is the original notional of the tranche, which is the cap of the liability held by the seller of the protection. Additionally, the single tranche CDO contract defines a schedule of accrual and payment dates, a fixed annualized periodic rate S, and sometimes also an upfront payment to be made by the buyer of the protection.

The cashflows are driven by the slice of the loss of the reference portfolio within the bounds $[k, K]$. The total loss sustained by the tranche since its inception until time T is

$$L_{CDO}(T) = (L(T) - k)^+ - (L(T) - K)^+ \tag{18}$$

(by definition $(x)^+ = x$ if $x > 0$ and $(x)^+ = 0$ otherwise). As soon as a positive jump ΔL_{CDO} in the quantity L_{CDO} is reported, the seller of the protection must pay the value ΔL_{CDO}

to the buyer of the protection. This is the only source of payments made by the seller of the protection. The payments made by the buyer of the protection are determined by the outstanding notional of the tranche $A_{CDO}(T)$ as a function of time T. The initial notional of the tranche is $A_{CDO}(0) = K - k$. As time goes by the notional of the tranche can get reduced "from below" because of losses but also possibly "from above" because of obligatory tranche amortization from recovery proceeds. The tranche gets touched by the amortization if the recovered proceeds exceed $A - K$ and can be early amortized up to its total original width $K - k$. The cumulative amortization amount applied to the tranche is

$$R_{CDO}(T) = (R(T) - (A - K))^+ - (R(T) - (A - k))^+. \tag{19}$$

Thus we can express the outstanding notional of the tranche as $A_{CDO}(T) = A_{CDO}(0) - L_{CDO}(T) - R_{CDO}(T)$. (We are safe subtracting the reductions due to loss and recovery independently because the part of the tranche reduced due to loss can never be subject to amortization.)

The outstanding notional of the tranche is monitored every day of each accrual period, and the fee is accrued on the outstanding notional of the tranche with the rate equal to the tranche spread S. The total accrued fee is paid on the payment date of the accrual period by the buyer of the protection to the seller of the protection. Thus if a default happened and reduced the notional of the tranche in the middle of an accrual period, the fee accrued since the beginning of the period while the asset was alive is included in the payment made by the protection buyer.

Let the accrual periods be $[0, T_1]$, $[T_1, T_2]$,... $[T_{k-1}, T_k]$. The analytical core for the pricing is the evaluation of the expectations

$$E_{\text{loss}}(T, k, K) = E[(L(T) - k)^+ - (L(T) - K)^+], \tag{20}$$

$$E_{\text{rec}}(T, A - K, A - k) = E[(R(T) - (A - K))^+ - (R(T) - (A - k))^+]. \tag{21}$$

Introducing the risk-free discount curve $D(t)$, the leg of the payments made by the protection seller (protection leg) is evaluated as

$$P_{\text{prot}} = \sum_i \left(E_{\text{loss}}(T_i, k, K) - E_{\text{loss}}(T_{i-1}, k, K)\right) D(T_i). \tag{22}$$

(We ignored here the exact timings of defaults, which is a commonly accepted approximation.) The leg of the payments made by the protection buyer (fee leg) is evaluated as

$$P_{\text{fee}} = \sum_i S \cdot \tau(T_{i-1}, T_i)(\alpha A_{CDO}(T_{i-1}) + \beta A_{CDO}(T_i))D(T_i). \tag{23}$$

Here $\tau(T_{i-1}, T_i)$ is the daycount fraction from T_{i-1} to T_i, and α, $\beta = 1 - \alpha$ are weights introduced to take into account the contribution of the accrued interest without introducing much extra complexity. When we set $\alpha = \beta = 0.5$ we mimic the effect of the accrued interest by effectively assuming that the defaults on the average happen in the middle of the accrual period. By setting $\alpha = 0$, $\beta = 1$ the accrued interest is excluded. The outstanding notional of the tranche is computed as $A_{CDO}(T_i) = (K - k) - E_{\text{loss}}(T_i, k, K) - E_{\text{rec}}(T_i, A - K, A - k)$, which completes the specification for the calculation of the fee leg. The present value of the tranche is the difference of the legs, $P_{\text{prot}} - P_{\text{fee}}$ for protection buyer, and $P_{\text{fee}} - P_{\text{prot}}$ for protection seller. Par spread of the single tranche CDO is the value of S that makes its present value equal to zero.

3.3 Deficiency of models with single correlation number and base correlations

Single tranche CDOs are actively tradable instruments, normally quoted simultaneously in terms of par spreads and implied correlations of the Gaussian copula model consistent with the par spreads. One of early disappointments with the Gaussian copula model was a counterintuitive behavior of implied correlations for mezzanine tranches. While for equity and senior tranches the solution for the implied correlation is normally unique, there usually is a mezzanine tranche for which the solution bifurcates into two branches with anomalously low and anomalously high correlation. Without going into further details of this widely publicized phenomenon here, we just give a simple argument why such a pathological behavior of the implied correlation should be expected of any reasonable model which only employs a single correlation number.

Indeed, it is easy to convince oneself that the par spread of an equity tranche should decrease with the increasing correlation while the par spread of a senior tranche should increase. Consider now any continuous deformation of the equity tranche into a senior tranche. For example, the tranche from 0 to 3% can be deformed into the tranche from 30% to 100% using the family of tranches $[u \cdot 30\%, 3\% + u \cdot 97\%]$ where u varies from 0 to 1. The dependence $S_u(\rho)$ of the par spread on the model correlation is a function of ρ which is decreasing for $u = 0$ and increasing for $u = 1$. Inevitably at some intermediate value $u = u_0$ the values $S_{u_0}(\rho = 0)$ and $S_{u_0}(\rho = 100\%)$ must cross. A highly degenerate situation is possible when $S_{u_0}(\rho)$ is completely independent of the correlation. However, this degeneracy is breakable by choosing a slightly different deformation of the tranches. In the generic case, the dependence $S_{u_0}(\rho)$ will have one or several bumps or dips, giving rise to bifurcations of the solution for implied correlations for almost any achievable value of spread. Thus any model that strives to solve inadequacies of the standard Gaussian copula must go beyond the single correlation number for every tranche.

Regardless of other arguments in favor and against, the base correlations introduced by McGinty et al. (2004) is a remedy against this specific problem because two rather than one correlation numbers are used to price a tranche. The expectations (20–21) are differences of independent stop-loss expectations corresponding to lower and upper bounds. It is possible to introduce a dependence of model parameters on the value of the bound and compute each stop-loss using the corresponding parameters. The base correlations approach consists in introducing a bound-dependent correlation $\rho = \rho(k)$ and computing $E_{\text{loss}}(T, 0, k)$ and $E_{\text{rec}}(T, 0, A - k)$ using $\rho(k)$, while computing $E_{\text{loss}}(T, 0, K)$ and $E_{\text{rec}}(T, 0, A - K)$ using $\rho(K)$, then obtaining $E_{\text{loss}}(T, k, K)$ as the difference $E_{\text{loss}}(T, 0, k) - E_{\text{loss}}(T, 0, K)$, and similarly for $E_{\text{rec}}(T, A - K, A - k)$. So far the calibration to market quotes has resulted in the base correlations that are free from the bifurcation phenomenon (although there are no theoretical grounds to claim that the base correlations will always remain well behaved under changing market conditions).

3.4 CDO^2

CDO^2 is a single tranche CDO based on a portfolio of CDO tranches. Let M be the number of underlying single tranche CDOs, and k_p, K_p, and A_p the attachment points and the total size of the underlying portfolio for the tranche p. Typically, M is from 2 to 12, each portfolio is based on 100 or 125 names, the total number of underlying names is 200–500, and there is significant overlap of underlying portfolios. The total amount at risk for the portfolio of tranches is $\mathcal{A} = \sum_p (K_p - k_p)$. The CDO^2 tranche has itself a pair of attachment points k and K. The original notional $A_{CDO^2}(t=0)$ of the tranche is equal to $K-k$, and the total loss sustained by the time T is

$$L_{CDO^2}(T) = (\mathcal{L}(T) - k)^+ - (\mathcal{L}(T) - K)^+ \tag{24}$$

where

$$\mathcal{L}(T) = \sum_p \left((L_p(T) - k_p)^+ - (L_p(T) - K_p)^+ \right) \tag{25}$$

is the total loss suffered by all the tranches and computed by restricting the loss $L_p(T)$ of each portfolio to the bounds of the corresponding tranche. The quantity $L_{CDO^2}(T)$ plays the same role in the definition of the cashflows of CDO^2 as the quantity $L_{CDO}(T)$ does in the definition of the cashflows of a single tranche CDO. Often CDO^{2}s are based on junior tranches and have vanishing probability of getting amortized, however for completeness we should also take recovery into account. Total amortization applied to the CDO^2 tranche is

$$R_{CDO^2}(T) = (\mathcal{R}(T) - (\mathcal{A} - K))^+ - (\mathcal{R}(T) - (\mathcal{A} - k))^+, \tag{26}$$

where

$$\mathcal{R}(T) = \sum_p \left((R_p(T) - (A_p - K_p))^+ - (R_p(T) - (A_p - k_p))^+ \right) \tag{27}$$

with $R_p(T)$ being total recovery in the underlying portfolio with index p. The rest of the definition of CDO^2 cashflows is the same as that for the single tranche CDO. The analytical core of the pricing is again the calculation of the expectations of bounded loss and recovery.

4 Analytical expansions I: single tranche CDO and single name portfolio

We resume the path of analysis of the models of correlated defaults with specific applications to single tranche CDOs and risk measures of single name portfolios. CDO^2 and portfolios with CDO tranches will be considered in the next section.

The general idea behind most analytical techniques is to expand or otherwise approximate the conditional CGF (11) by a more tractable expression. This is the same idea used in a typical derivation of central limit theorems in probability theory. Expanding the conditional CGF up to the second order around $\xi = 0$ we get the Gaussian normal approximation for the distribution of loss, matching the first two conditional moments of the exact distribution. The explicit expressions for the moments read

$$\Lambda(\mathbf{X}) = \sum_a w_a \mu_a(\mathbf{X}), \tag{28}$$

$$M_2(\mathbf{X}) = \sum_a w_a^2 \mu_a(\mathbf{X})(1 - \mu_a(\mathbf{X})). \tag{29}$$

The advantage of going through the expression for the CGF rather than just citing the central limit theorem and postulating the proxy normal distribution as done by Shelton (2004) is that it is possible in principle to evaluate corrections by retaining higher terms in the expansion of the CGF. A particularly convenient way of doing so consists in keeping only the quadratic terms in the exponent and expanding the exponential of cubic and further terms,

$$\exp(\mathcal{K}(\xi|\mathbf{X})) = e^{\xi\Lambda(\mathbf{X})+\frac{1}{2}\xi^2 M_2(\mathbf{X})}\left(1+\xi^3 a_3(\mathbf{X})+\xi^4 a_4(\mathbf{X})+\ldots\right), \tag{30}$$

so that the calculation of the inverse Laplace transform (12) only involves Gaussian integration. This is the Edgeworth expansion. The calculation of higher terms in the Edgeworth expansion is simple but is of limited practical value because the series is asymptotic and does not converge to the exact answer.

It turns out that $\xi = 0$ is not the optimal point to expand the CGF. A significantly better approximation is achieved if the expansion to the quadratic terms is made around the stationary point of the exponential, the saddlepoint.

4.1 Portfolio risk measures

The idea of the saddlepoint method (pioneered by Daneils (1954) in applications to statistics and by Martin et al. (2001b) for credit risk measures) is to identify the leading contribution to the integral (12) as coming from the vicinity of the real stationary point of the exponent, that is from the root $\xi = \xi_0$ of the equation $\mathcal{K}'(\xi|\mathbf{X}) - L = 0$, with

$$\mathcal{K}'(\xi|\mathbf{X}) = \sum_a \frac{w_a \mu_a \exp(\xi w_a)}{1-\mu_a+\mu_a\exp(\xi w_a)}. \tag{31}$$

(The dependence of μ_a and also of the saddlepoint ξ_0 on the central factor $\mathbf{X}$ in this and subsequent equations is understood.) The leading order for the conditional probability of loss is

$$P_{s0}(L|\mathbf{X}) = \frac{e^{\mathcal{K}(\xi_0|\mathbf{X})-\xi_0 L}}{\sqrt{2\pi\mathcal{K}''(\xi_0|\mathbf{X})}}, \tag{32}$$

where

$$\mathcal{K}''(\xi|\mathbf{X}) = \sum_a \frac{w_a^2\mu_a(1-\mu_a)\exp(\xi w_a)}{(1-\mu_a+\mu_a\exp(\xi w_a))^2}. \tag{33}$$

Corrections to the leading saddlepoint approximation are generated by the higher order terms in the expansion of the exponential of the CGF around the saddlepoint,

$$\exp(\mathcal{K}(\xi|\mathbf{X})-\xi L) = e^{\mathcal{K}(\xi_0|\mathbf{X})-\xi_0 L+\frac{1}{2}\mathcal{K}''(\xi_0|\mathbf{X})(\xi-\xi_0)^2}\left(1+\tfrac{1}{6}\mathcal{K}'''(\xi_0|\mathbf{X})(\xi-\xi_0)^3+\ldots\right). \tag{34}$$

It can be shown that the contribution of the term $(\xi_-\xi_0)^s\mathcal{K}^{(s)}$ scales as $N^{1-s/2}$ with increasing number N of underlying names (we always assume that the weights w_a are of comparable magnitude so that no one asset or small subgroup of assets dominates the credit exposure). The correction to the probability density $P(L|\mathbf{X})$ of the order $N^{-1/2}$ from the term with $(\xi-\xi_0)^3$ vanishes. The first non-vanishing correction $P_{s1}(L|\mathbf{X})$ is of the order N^{-1} and comes from two higher order terms in the expansion (34), namely $(1/24)\mathcal{K}^{\mathrm{IV}}(\xi_0|\mathbf{X})(\xi-\xi_0)^4$ and $(1/72)(\mathcal{K}'''(\xi_0|\mathbf{X}))^2(\xi-\xi_0)^6$. The well-known result of Daniels (1954) for the correction reads

$$P_{s1}(L|\mathbf{X}) = \frac{e^{\mathcal{K}(\xi_0|\mathbf{X})-\xi_0 L}}{\sqrt{2\pi\mathcal{K}''(\xi_0|\mathbf{X})}}\left(\frac{1}{8}\frac{\mathcal{K}^{\mathrm{IV}}(\xi_0|\mathbf{X})}{(\mathcal{K}''(\xi_0|\mathbf{X}))^2}-\frac{5}{24}\frac{(\mathcal{K}'''(\xi_0|\mathbf{X}))^2}{\mathcal{K}''(\xi_0|\mathbf{X}))^3}\right). \tag{35}$$

The calculation of the asymptotic series of higher order corrections coming from the expansion (34) will not result in infinite accuracy improvement. In particular it does not allow recovery of the discrete nature of the distribution of loss. (The exact discrete distribution can be recovered if the contributions of other saddlepoints off the real axis in the complex plane are accounted for). Practically it is not necessary to go beyond the leading saddlepoint approximation and the first correction to achieve sufficient accuracy.

With application to credit portfolio VaR and expected shortfall we should apply the saddlepoint method directly to the quantities of interest without computing the probability density of loss. For the probability of loss exceeding the threshold K we get

$$P[L \geq K|\mathbf{X}] = \frac{1}{2\pi i}\int_K^\infty dL \int_{c-i\infty}^{c+i\infty} \exp(\mathcal{K}(\xi|\mathbf{X}) - \xi L) d\xi. \tag{36}$$

Integrating over L first we obtain the expression

$$P[L \geq K|\mathbf{X}] = \frac{1}{2\pi i}\int_{c-i\infty}^{c+i\infty} \frac{\exp(\mathcal{K}(\xi|\mathbf{X}) - \xi K)}{\xi} d\xi, \tag{37}$$

which is ready for the application of the saddlepoint method. Similarly for the shortfall integral

$$E[L|L \geq K;\mathbf{X}] = (P[L \geq K])^{-1}\frac{1}{2\pi i}\int_K^\infty dL \int_{c-i\infty}^{c+i\infty} \exp(\mathcal{K}(\xi|\mathbf{X}) - \xi L) L d\xi \tag{38}$$

we get the expression

$$E[L|L \geq K;\mathbf{X}] = (P[L \geq K])^{-1}\left(\frac{1}{2\pi i}\int_{c-i\infty}^{c+i\infty} \frac{\exp(\mathcal{K}(\xi|\mathbf{X}) - \xi K)}{\xi^2} d\xi + KP[L \geq K|\mathbf{X}]\right). \tag{39}$$

For the particular value of strike K equal to expected loss, $\Lambda(\mathbf{X})$, the saddlepoint is $\xi_0 = 0$, and $\mathcal{K}(\xi_0|\mathbf{X}) = 0$, $\mathcal{K}''(\xi_0|\mathbf{X}) = M_2(\mathbf{X})$. For $K < \Lambda(\mathbf{X})$ the saddlepoint is negative, $\xi_0 < 0$, and for $K > \Lambda(\mathbf{X})$ it is positive, $\xi_0 > 0$. In the case of a negative saddlepoint, care has to be taken in deforming the contour of integration across the pole at origin. The contribution from the residue in the pole should be computed exactly before expanding near the saddlepoint.

After taking the average over the central factor $\mathbf{X}$, the leading saddlepoint contribution to the tail probability is

$$P[L \geq K] = E_{\mathbf{X}}\left[\theta(-\xi_0) + e^{\mathcal{K}(\xi_0|\mathbf{X}) - \xi_0 K} J_1(\mathcal{K}''(\xi_0|\mathbf{X}), \xi_0)\right]. \tag{40}$$

The leading saddlepoint contribution to the expected shortfall is

$$E[L|L \geq K] = \frac{E_{\mathbf{X}}\left[\theta(-\xi_0)(\sum_a w_a \mu_a - K) + e^{\mathcal{K}(\xi_0|\mathbf{X}) - \xi_0 K} J_2(\mathcal{K}''(\xi_0|\mathbf{X}), \xi_0)\right]}{P(L \geq K)} + K, \tag{41}$$

where we introduced the notation $J_s(m, \xi_0)$ for the integral

$$J_s(m, \xi_0) = \frac{1}{2\pi i}\int_{\xi_0 - i\infty}^{\xi_0 + i\infty} \exp\left(\tfrac{1}{2}m(\xi - \xi_0)^2\right) d\xi/\xi^s. \tag{42}$$

The values of this integral with $s = 0, 1, 2$ are as follows

$$J_0(m, \xi_0) = 1/\sqrt{2\pi m}, \tag{43}$$

$$J_1(m, \xi_0) = \mathrm{sign}(\xi_0)\exp(\tfrac{1}{2}m\xi_0^2)\mathcal{N}(-\sqrt{m}|\xi_0|), \tag{44}$$

$$J_2(m, \xi_0) = \sqrt{m/2\pi} - m|\xi_0|\exp(\tfrac{1}{2}m\xi_0^2)\mathcal{N}(-\sqrt{m}|\xi_0|). \tag{45}$$

Note a useful relationship $J_2(m, \xi_0) = m(J_0(m, \xi_0) - \xi_0 J_1(m, \xi_0))$. In all subsequent expressions these integrals are used with the parameters $m = \mathcal{K}''(\xi_0|\mathbf{X})$ and ξ_0, which will sometimes be omitted.

The lowest correction to the tail probability and expected shortfall comes from the cubic term in Eq. (34).

$$\Delta P[L \geq K] = E_{\mathbf{X}} \left[\frac{1}{6} \mathcal{K}'''(\xi_0|\mathbf{X}) e^{\mathcal{K}(\xi_0|\mathbf{X}) - \xi_0 K} \left(\left(\xi_0^2 - \frac{1}{\mathcal{K}''(\xi_0|\mathbf{X})} \right) J_0 - \xi_0^3 J_1 \right) \right]. \tag{46}$$

The calculation of the VaR requires computing the strike K to match the given probability in the tail, $1 - \alpha$. This can be done using a one-dimensional iterative solver. (Note that this should be done using the unconditional probability averaged over the global factor $\mathbf{X}$. Even though it is possible to define VaR conditional on the central factor, the unconditional VaR cannot be restored from it by integration.)

The calculation of VaR decomposition into marginal contributions from individual assets can be done without multiple repetitions of solving the equation $P[L \geq K] = 1 - \alpha$, as was shown by Gourieroux et al. (2000) and Martin et al. (2001c). We do it using the following identity for partial derivatives

$$\frac{\partial \mathrm{VaR}_\alpha(L)}{\partial w_a} = - \frac{\frac{\partial P[L \geq K]}{\partial w_a} |_{K = \mathrm{VaR}_\alpha(L)}}{\frac{\partial P[L \geq K]}{\partial K} |_{K = \mathrm{VaR}_\alpha(L)}}. \tag{47}$$

The quantity in the denominator is the probability density of loss taken at the VaR level, and the quantity in the numerator is the sensitivity of the tail probability to the asset's weight. This can be written as

$$\frac{\partial \mathrm{VaR}_\alpha(L)}{\partial w_a} = \left. \frac{E_{\mathbf{X}} \left[\int \frac{\mu_a \exp(\xi w_a)}{1 - \mu_a + \mu_a \exp(\xi w_a)} \exp(\mathcal{K}(\xi|\mathbf{X}) - \xi K) d\xi \right]}{E_{\mathbf{X}} \left[\int \exp(\mathcal{K}(\xi|\mathbf{X}) - \xi K) d\xi \right]} \right|_{K = \mathrm{VaR}_\alpha(L)}. \tag{48}$$

The integrals over ξ can be taken in the leading saddlepoint approximation. The saddlepoint ξ_0 is determined by the equation $\mathcal{K}'(\xi_0|\mathbf{X}) = \mathrm{VaR}_\alpha(L)$. We obtain in this way the following decomposition

$$\mathrm{VaR}_\alpha(L) = \sum_a \left. \frac{E_{\mathbf{X}} \left[P_{s0}(K|\mathbf{X}) \frac{w_a \mu_a \exp(\xi_0 w_a)}{1 - \mu_a + \mu_a \exp(\xi_0 w_a)} \right]}{E_{\mathbf{X}}[P_{s0}(K|\mathbf{X})]} \right|_{K = \mathrm{VaR}_\alpha(L)}. \tag{49}$$

Note that despite the approximation the equality (49) holds exactly because of the defining equation of the saddlepoint, $\mathcal{K}'(\xi_0|\mathbf{X}) = \mathrm{VaR}_\alpha(L)$ (see Eq. (31)). Thus we obtain the decomposition of VaR with minimal computational cost.

The calculation of expected shortfall starts with the calculation of VaR for the requested confidence level. This value of VaR is then plugged for K in Eq. (41). The probability in the denominator is exactly $1 - \alpha$. The final expression in the leading saddlepoint approximation is

$$\mathrm{ES}_\alpha(L) = \left. \frac{E_{\mathbf{X}} \left[\theta(-\xi_0) (\sum_a w_a \mu_a - K) + e^{\mathcal{K}(\xi_0|\mathbf{X}) - \xi_0 K} J_2(\mathcal{K}''(\xi_0|\mathbf{X}), \xi_0) \right]}{1 - \alpha} + K \right|_{K = \mathrm{VaR}_\alpha(L)}. \tag{50}$$

The first correction to this expression comes from the correction to the stop-loss, written down in the next section.

To obtain the shortfall decomposition, we start with the exact expression

$$\frac{\partial ES_\alpha(L)}{\partial w_a} = \frac{E_{\mathbf{X}}\left[\frac{1}{2\pi i}\int_{c-i\infty}^{c+i\infty}\frac{\exp(\mathcal{K}(\xi|\mathbf{X})-\xi K)}{\xi}\frac{\mu_a\exp(\xi w_a)}{1-\mu_a+\mu_a\exp(\xi w_a)}d\xi\right]_{K=\mathrm{VaR}_\alpha(L)}}{1-\alpha}. \tag{51}$$

To get a decomposition that is consistent with the leading saddlepoint approximation, it is necessary to keep the linear term in the expansion of the fraction $\mu_a\exp(\xi w_a)/(1-\mu_a+\mu_a\exp(\xi w_a))$ around the saddlepoint ξ_0,

$$\frac{\mu_a\exp(\xi w_a)}{1-\mu_a+\mu_a\exp(\xi w_a)} = \frac{\mu_a\exp(\xi_0 w_a)}{1-\mu_a+\mu_a\exp(\xi_0 w_a)} + \frac{w_a\mu_a(1-\mu_a)\exp(\xi_0 w_a)}{(1-\mu_a+\mu_a\exp(\xi_0 w_a))^2}(\xi-\xi_0)+\dots \tag{52}$$

Before applying the expansion we need to deform the contour of integration so that it passes through the saddlepoint, which produces a contribution from the residue at origin if the saddlepoint is negative. The first term in the expansion (52) generates the decomposition of the addendum equal to VaR in the expression (50) for expected shortfall. The integral over ξ cancels against $1-\alpha$ in the denominator (in order for this cancellation to work in the case $\xi_0 < 0$, we have to pass the contour back through the origin, picking up another residue term). The second term in (52) generates the term with J_2, where the integral over ξ taken in the leading saddlepoint approximation to ensure consistency of the decomposition. The final decomposition reads

$$\begin{aligned} ES_\alpha(L) &= \sum_a E_{\mathbf{X}}\left[\frac{w_a\mu_a\exp(\xi_0 w_a)}{1-\mu_a+\mu_a\exp(\xi_0 w_a)} + \theta(-\xi_0)\frac{w_a\mu_a - \frac{w_a\mu_a\exp(\xi_0 w_a)}{1-\mu_a+\mu_a\exp(\xi_0 w_a)}}{1-\alpha}\right] \\ &+ \sum_a \frac{E_{\mathbf{X}}\left[e^{\mathcal{K}(\xi_0|\mathbf{X})-\xi_0 K}J_2(\mathcal{K}''(\xi_0|\mathbf{X}),\xi_0)\frac{w_a^2\mu_a(1-\mu_a)\exp(\xi w_a)}{(1-\mu_a+\mu_a\exp(\xi w_a))^2\mathcal{K}''(\xi_0|\mathbf{X})}\right]_{K=\mathrm{VaR}_\alpha(L)}}{1-\alpha}. \end{aligned} \tag{53}$$

The Eqs. (31) and (33) neatly conspire to ensure the equality.

4.2 Single tranche CDO

For single tranche CDO pricing we should apply the saddlepoint method directly to the stop-loss integral conditional on the global factor

$$E[(L-K)^+|\mathbf{X}] = \frac{1}{2\pi i}\int_K^\infty dL\int_{c-i\infty}^{c+i\infty}\exp(\mathcal{K}(\xi|\mathbf{X})-\xi L)(L-K)d\xi. \tag{54}$$

Integrating over L first we obtain the expression

$$E[(L-K)^+|\mathbf{X}] = \frac{1}{2\pi i}\int_{c-i\infty}^{c+i\infty}\frac{\exp(\mathcal{K}(\xi|\mathbf{X})-\xi K)}{\xi^2}d\xi. \tag{55}$$

Approximating the exponential of $\mathcal{K}(\xi)-\xi K$ around the saddle point we retain terms up to the quadratic in the exponent and expand the rest

$$\exp(\mathcal{K}(\xi)-\xi K) = \exp\left(\mathcal{K}(\xi_0)-\xi_0 K+\frac{1}{2}\mathcal{K}''(\xi_0)(\xi-\xi_0)^2\right)\left(1+\frac{1}{6}\mathcal{K}'''(\xi_0)(\xi-\xi_0)^3+\dots\right). \tag{56}$$

Evaluation of the integrals yields the leading saddlepoint contribution to the conditional stop-loss (with the usual care taken of the residue term)

$$E[(L-K)^+] = E_{\mathbf{X}}\left[\theta(-\xi_0)(\sum_a w_a\mu_a - K) + e^{\mathcal{K}(\xi_0|\mathbf{X})-\xi_0 K} J_2(\mathcal{K}''(\xi_0|\mathbf{X}),\xi_0)\right] \tag{57}$$

and the first correction

$$\Delta E[(L-K)^+] = E_{\mathbf{X}}\left[\tfrac{1}{6}\xi_0\mathcal{K}'''(\xi_0|\mathbf{X})e^{\mathcal{K}(\xi_0|\mathbf{X})-\xi_0 K}\left(-2J_0 + 3\xi_0 J_1 - \xi_0^2 J_2\right)\right]. \tag{58}$$

4.3 Granularity adjustment

We note here that Gordy's granularity adjustment, considered in Gordy (2003), Wilde (2001), Martin and Wilde (2002), and Emmer and Tasche (2005) in the context of portfolio risk measures can also be regarded as a variant of the CGF expansion technique. To derive the granularity adjustment we again expand the CGF but this time leave only the linear term in the exponent

$$\exp(\mathcal{K}(\xi|\mathbf{X})) = e^{\xi\Lambda(\mathbf{X})}(1 + \tfrac{1}{2}\xi^2 M_2(\mathbf{X}) + \ldots). \tag{59}$$

The contour integral of the term with $M_2(\mathbf{X})$ (and also of all higher terms) is singular, but the singularity is removed by the residual integration over the central factor $\mathbf{X}$. Specifically, we get for the tail probability

$$P[L \geq K] = E_{\mathbf{X}}[\theta(\Lambda(\mathbf{X}) - K) + \tfrac{1}{2}M_2(\mathbf{X})\delta'(\Lambda(\mathbf{X}) - K) + \ldots] \tag{60}$$

and for the stop-loss option

$$E[(L-K)^+] = E_{\mathbf{X}}[(\Lambda(\mathbf{X}) - K)^+ + \tfrac{1}{2}M_2(\mathbf{X})\delta(\Lambda(\mathbf{X}) - K) + \ldots]. \tag{61}$$

Dirac's delta-function, $\delta(x)$, is defined as an integration kernel by the property

$$\int_{-\infty}^{\infty} dx\, \delta(x)\, f(x) = f(0)$$

for any function $f(x)$. Higher derivatives also have a meaning as integration kernels and respect the rule of integration by parts,

$$\int_{-\infty}^{\infty} dx\, \delta^{(n)}(x)\, f(x) = (-1)^n\, f^{(n)}(x)|_{x=0}.$$

For a multi-dimensional central factor $\mathbf{X}$ the singular integrals reduce to the surface $\Lambda(\mathbf{X}) = K$ and can be taken numerically. For a one-dimensional central factor X with the probability density $p(x)$ (which can be any continuous probability density, not necessarily normal) the singular integrals are fully determined by the root x_0 where average loss equals strike, $\Lambda(x_0) = K$, and can be written down explicitly,

$$P[L \geq K] \approx \int dx\, p(x)\theta(\Lambda(x) - K) - \frac{1}{2|\Lambda'(x_0)|}\frac{d}{dx}\frac{M_2(x)p(x)}{\Lambda'(x_0)}\Big|_{x=x_0}, \tag{62}$$

$$E[(L-K)^+] \approx \int dx\, p(x)(\Lambda(x) - K)^+ + \frac{M_2(x_0)p(x_0)}{2|\Lambda'(x_0)|}. \tag{63}$$

In the case of several roots, a summation over all the roots is needed. If there is no root, the granularity adjustment is absent. Eq. (62) is the granularity adjustment to the tail probability, different in form but equivalent to the standard expression for the granularity adjustment to VaR given by Gordy (2003) and Wilde (2001). Eq. (63) is its counterpart for the stop-loss option. Evaluation of higher corrections due to higher moments in the CGF expansion is straightforward and can be implemented easily. However, as with the Edgeworth expansion, the calculation of higher corrections has limited practical value because the lack of analyticity precludes the asymptotic series expansion from convergence to the exact value. (The discrepancy between the saddlepoint approach and granularity adjustment noted in Martin and Wilde (2002) is another manifestation of this non-analyticity.) As we will see shortly the granularity correction usually improves the result, however the method remains considerably less accurate than the conditional normal approximation for portfolio loss and much less accurate than the saddlepoint.

4.4 Numerical results

In this section we compare numerical results obtained for the stop-loss using the various analytical approximation. We consider as an example a portfolio of 125 assets with loss weights drawn from a uniform distribution between 0.5 and 0.7. In the first set of runs the average credit spread on the horizon of the calculation is 100.31 bp, corresponding to the average default probability 1.65%. The value of the total expected loss is $E[L] = 1.2267$ and the maximum possible loss is $L_{max} = 74.1341$. We fix a set of upper bounds K in terms of certain percentages of L_{max}. The reported quantity is $(E[L] - E[(L-K)^+])/E[L]$ which is the expectation of loss bounded from 0 to the upper strike and normalized to the expected loss. The calculations are done assuming the standard Gaussian copula model with a one-dimensional central factor for several values of the correlation.

The baseline for comparison is the result from a slow convolution calculation with adaptively chosen grid (a simple recursion scheme is not sufficient because of incommensurability of loss weights).

Adaptive grid convolution, $\langle p_d \rangle = 1.65\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	0.498249	0.804019	0.941830	0.997456	1.000000	1.000000	1.000000
10%	0.417109	0.669786	0.815121	0.940278	0.995735	0.999658	1.000000
20%	0.353641	0.567724	0.703215	0.849748	0.965540	0.990675	0.999749
30%	0.299737	0.482701	0.605586	0.755074	0.910081	0.962535	0.996520
40%	0.252314	0.408583	0.518303	0.662446	0.838825	0.915200	0.985135
50%	0.209781	0.342266	0.438783	0.573092	0.757836	0.852012	0.961210

The large portfolio model is the first term in Eq. (63) without the granularity adjustment. It is quite a crude approximation as can be seen from the following results.

Large portfolio, $\langle p_d \rangle = 1.65\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	0.604318	1.000000	1.000000	1.000000	1.000000	1.000000	1.000000
10%	0.507550	0.764342	0.885192	0.969137	0.998415	0.999896	1.000000
20%	0.421121	0.631890	0.753850	0.881343	0.974069	0.993394	0.999833
30%	0.352210	0.530896	0.643900	0.781356	0.921492	0.967945	0.997163
40%	0.292355	0.443801	0.548219	0.685004	0.849707	0.922671	0.986479
50%	0.239592	0.371070	0.465097	0.591096	0.769500	0.858077	0.963876

Large portfolio error, $\langle p_d \rangle = 1.65\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	0.106069	0.195981	0.058170	0.002544	0.000000	0.000000	-0.000000
10%	0.090442	0.094556	0.070071	0.028858	0.002680	0.000238	-0.000000
20%	0.067481	0.064166	0.050635	0.031595	0.008529	0.002718	0.000084
30%	0.052473	0.048195	0.038315	0.026283	0.011411	0.005410	0.000643
40%	0.040041	0.035218	0.029916	0.022558	0.010881	0.007471	0.001345
50%	0.029811	0.028804	0.026313	0.018004	0.011664	0.006064	0.002666

The next pair of tables shows that the granularity adjustment gives only a modest improvement in the majority of cases. Note that the granularity adjustment is not applicable without correlation, $\rho = 0$, because there is no dependence on the central factor $\mathbf{X}$ and no solution to the equation $\Lambda(x_0) = K$.

Large portfolio with granularity adjustment, $\langle p_d \rangle = 1.65\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	0.604318	1.000000	1.000000	1.000000	1.000000	1.000000	1.000000
10%	0.443194	0.702935	0.844178	0.954317	0.997338	0.999807	1.000000
20%	0.376507	0.592572	0.723167	0.863614	0.969336	0.991995	0.999787
30%	0.319337	0.502536	0.620512	0.765482	0.914906	0.964955	0.996820
40%	0.267561	0.422544	0.530148	0.671659	0.842762	0.918746	0.985636
50%	0.220866	0.355012	0.451160	0.580232	0.762986	0.853861	0.962543

Large portfolio with granularity adjustment error, $\langle p_d \rangle = 1.65\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	0.106069	0.195981	0.058170	0.002544	0.000000	0.000000	-0.000000
10%	0.026085	0.033149	0.029056	0.014039	0.001603	0.000149	-0.000000
20%	0.022867	0.024849	0.019952	0.013866	0.003797	0.001319	0.000037
30%	0.019600	0.019835	0.014927	0.010409	0.004825	0.002420	0.000301
40%	0.015247	0.013961	0.011845	0.009213	0.003937	0.003546	0.000502
50%	0.011085	0.012747	0.012377	0.007140	0.005150	0.001848	0.001333

However it cannot compete with the normal proxy approach, in which the stop-loss is given by the expression from Shelton (2004)

$$E[(L-K)^+] = E_{\mathbf{X}}\Bigg[(\Lambda(\mathbf{X})-K)\mathcal{N}\left(\frac{\Lambda(\mathbf{X})-K}{\sqrt{M_2(\mathbf{X})}}\right) + \sqrt{\frac{M_2(\mathbf{X})}{2\pi}}\exp\left(-\frac{(\Lambda(\mathbf{X})-K)^2}{2M_2(\mathbf{X})}\right)\Bigg]. \quad (64)$$

Normal proxy, $\langle p_d \rangle = 1.65\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	0.482537	0.816231	0.959354	0.999658	1.000000	1.000000	1.000000
10%	0.413002	0.675863	0.820795	0.943272	0.996052	0.999662	1.000000
20%	0.352786	0.571816	0.706292	0.851585	0.966057	0.990824	0.999745
30%	0.299991	0.485671	0.607544	0.756253	0.910546	0.962727	0.996546
40%	0.252934	0.410792	0.519632	0.663231	0.839190	0.915378	0.985170
50%	0.210464	0.343911	0.439706	0.573622	0.758109	0.852158	0.961246

Normal proxy error, $\langle p_d \rangle = 1.65\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	-0.015712	0.012212	0.017524	0.002202	0.000000	0.000000	-0.000000
10%	-0.004106	0.006077	0.005674	0.002994	0.000317	0.000004	-0.000000
20%	-0.000855	0.004092	0.003077	0.001838	0.000518	0.000148	-0.000004
30%	0.000254	0.002970	0.001958	0.001180	0.000465	0.000192	0.000027
40%	0.000620	0.002209	0.001329	0.000785	0.000365	0.000178	0.000035
50%	0.000683	0.001646	0.000922	0.000529	0.000272	0.000146	0.000036

The leading saddlepoint approximation (57) clearly works better than the normal proxy.

Saddlepoint, leading order, $\langle p_d \rangle = 1.65\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	0.511338	0.797625	0.936926	0.997185	1.000000	1.000000	1.000000
10%	0.418538	0.670274	0.813768	0.939394	0.995661	0.999626	1.000000
20%	0.353190	0.568288	0.702624	0.849237	0.965466	0.990664	0.999741
30%	0.298816	0.483119	0.605263	0.754709	0.909998	0.962515	0.996531
40%	0.251406	0.409013	0.518178	0.662152	0.838784	0.915170	0.985142
50%	0.208992	0.342539	0.438671	0.572845	0.757743	0.851984	0.961208

Saddlepoint, leading order error, $\langle p_d \rangle = 1.65\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	0.013089	-0.006393	-0.004904	-0.000271	-0.000000	0.000000	-0.000000
10%	0.001429	0.000488	-0.001353	-0.000884	-0.000074	-0.000032	-0.000000
20%	-0.000451	0.000564	-0.000591	-0.000511	-0.000073	-0.000011	-0.000008
30%	-0.000920	0.000418	-0.000322	-0.000365	-0.000083	-0.000020	0.000011
40%	-0.000908	0.000430	-0.000125	-0.000294	-0.000041	-0.000030	0.000007
50%	-0.000789	0.000273	-0.000112	-0.000247	-0.000094	-0.000029	-0.000002

Finally, the first correction (58) to the saddlepoint approximation can further improve the results. (As shown by Taras et al. (2005), the saddlepoint approximation with corrections works good for the tail probability as well.)

Saddlepoint with first correction, $\langle p_d \rangle = 1.65\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	0.500348	0.807992	0.942917	0.997396	1.000000	1.000000	1.000000
10%	0.417634	0.672211	0.815709	0.940271	0.995730	0.999631	1.000000
20%	0.353757	0.569447	0.703633	0.849773	0.965581	0.990695	0.999742
30%	0.299686	0.483993	0.605896	0.755099	0.910127	0.962562	0.996534
40%	0.252206	0.409579	0.518538	0.662464	0.838864	0.915222	0.985148
50%	0.209673	0.343027	0.438957	0.573106	0.757866	0.852028	0.961220

Saddlepoint with first correction error, $\langle p_d \rangle = 1.65\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	0.002099	0.003974	0.001088	-0.000060	-0.000000	0.000000	-0.000000
10%	0.000525	0.002425	0.000588	-0.000007	-0.000005	-0.000027	-0.000000
20%	0.000116	0.001723	0.000418	0.000026	0.000042	0.000019	-0.000007
30%	-0.000051	0.001292	0.000311	0.000026	0.000046	0.000028	0.000014
40%	-0.000108	0.000996	0.000235	0.000018	0.000039	0.000022	0.000014
50%	-0.000107	0.000762	0.000173	0.000013	0.000030	0.000016	0.000010

The repetition of the same runs with higher average default probability, 4.05%, which corresponds to average credit spread 250.79 bp and expected loss $E[L] = 3.0081$, leads to the same ranking of the methods. Below are the baseline for the quantity $(E[L]-E[(L-K)^+])/E[L]$ and the results for the deviation from the baseline.

Adaptive grid convolution, $\langle p_d \rangle = 4.05\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	0.243699	0.474085	0.672744	0.913892	0.999697	1.000000	1.000000
10%	0.225383	0.416175	0.569004	0.772490	0.957416	0.992311	0.999977
20%	0.203500	0.365358	0.493076	0.671730	0.882221	0.955519	0.997413
30%	0.181252	0.320065	0.429583	0.588393	0.802738	0.899871	0.985599
40%	0.159461	0.278835	0.373583	0.515080	0.724094	0.834693	0.961248
50%	0.138391	0.240716	0.322689	0.448314	0.647009	0.764064	0.924305

Large portfolio error, $\langle p_d \rangle = 4.05\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	0.002745	0.018803	0.066589	0.086108	0.000303	0.000000	0.000000
10%	0.017073	0.037674	0.046951	0.044331	0.015590	0.003603	0.000009
20%	0.021078	0.031181	0.034354	0.029817	0.016289	0.007818	0.000700
30%	0.020344	0.024795	0.025694	0.022656	0.014055	0.008294	0.001720
40%	0.017830	0.019859	0.018969	0.018921	0.011388	0.010113	0.002515
50%	0.015183	0.016351	0.015951	0.012765	0.009210	0.006892	0.002594

Large portfolio with granularity adjustment error, $\langle p_d \rangle = 4.05\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	0.002745	0.018803	0.066589	0.086108	0.000303	0.000000	0.000000
10%	0.009030	0.014445	0.016027	0.016111	0.007292	0.001869	-0.000000
20%	0.008116	0.011895	0.013510	0.011174	0.006673	0.003423	0.000344
30%	0.007536	0.009464	0.010202	0.008850	0.005423	0.003153	0.000693
40%	0.006514	0.007677	0.007045	0.008286	0.004037	0.005100	0.000933
50%	0.005671	0.006723	0.006694	0.004488	0.003110	0.002366	0.000701

Normal proxy error, $\langle p_d \rangle = 4.05\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	-0.004433	-0.006973	-0.006228	0.005193	0.000263	0.000000	0.000000
10%	-0.003870	-0.001656	-0.000506	0.001211	0.000892	0.000227	-0.000016
20%	-0.002359	-0.000114	0.000199	0.000676	0.000572	0.000271	0.000023
30%	-0.001300	0.000391	0.000336	0.000448	0.000380	0.000220	0.000046
40%	-0.000642	0.000545	0.000339	0.000316	0.000261	0.000164	0.000048
50%	-0.000249	0.000556	0.000301	0.000228	0.000182	0.000119	0.000042

Saddlepoint, leading order error, $\langle p_d \rangle = 4.05\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	0.000607	0.002707	0.004500	-0.004138	-0.000014	-0.000000	0.000000
10%	0.001478	0.001634	0.000367	-0.000403	-0.000186	-0.000050	-0.000018
20%	0.000796	0.000886	0.000080	-0.000202	-0.000095	-0.000020	0.000004
30%	0.000291	0.000585	0.000049	-0.000095	-0.000018	-0.000004	0.000015
40%	0.000045	0.000406	0.000006	-0.000123	-0.000039	-0.000013	0.000005
50%	-0.000065	0.000273	-0.000059	-0.000097	-0.000050	0.000029	-0.000001

Saddlepoint with first correction error, $\langle p_d \rangle = 4.05\%$							
$\rho \setminus K$	0.01	0.02	0.03	0.05	0.1	0.15	0.3
0%	0.000211	0.000811	0.000297	0.000069	0.000000	0.000000	0.000000
10%	0.000335	0.000924	0.000247	0.000039	0.000040	-0.000000	-0.000018
20%	0.000210	0.000792	0.000199	0.000032	0.000056	0.000030	0.000007
30%	0.000116	0.000661	0.000158	0.000021	0.000048	0.000030	0.000018
40%	0.000060	0.000548	0.000129	0.000012	0.000037	0.000022	0.000016
50%	0.000029	0.000449	0.000102	0.000012	0.000028	0.000016	0.000010

4.5 Saddlepoint method as normal proxy

We derived both the normal proxy method and the saddlepoint method starting from the same conditional CGF but expanding it around different points. Regardless of the expansion point, any quadratic CGF corresponds to a Gaussian normal distribution. For the normal proxy method the mean value $\Lambda(\mathbf{X})$ and the variance $M_2(\mathbf{X})$ are given by Eqs. (28-29). It is possible to give a description of the leading saddlepoint approximation in the same terms by applying the inverse Laplace transform to the expanded CGF (34). The moments of the resulting shifted Gaussian distribution depend on the location of the saddlepoint ξ_0,

$$\tilde{\Lambda}(\mathbf{X}) = K - \xi_0 \mathcal{K}''(\xi_0|\mathbf{X}), \quad (65)$$

$$\tilde{M}_2(\mathbf{X}) = \mathcal{K}''(\xi_0|\mathbf{X}). \quad (66)$$

The normal distribution of loss with the moments (65-66) gives better results because it is tuned to the specific problem, such as the calculation of stop-loss or tail probability with a given threshold K. The standard normal proxy moments (28) and (29) are restored only for K equal to expected conditional loss $\Lambda(\mathbf{X})$, in which case $\xi_0 = 0$.

Results from the shifted normal proxy with moments (65) and (66) are slightly different from the standard saddlepoint results because the normalization of the total probability to 1 is enforced. The standard saddlepoint approximations break the normalization as the value of the expanded CGF at $\xi = 0$ deviates from 0. This deviation can be eliminated by a replacement of the factor $\exp(\mathcal{K}(\xi_0|\mathbf{X}) - \xi_0 K))$ by $\exp(-\frac{1}{2}\mathcal{K}''(\xi_0|\mathbf{X}))$ in the leading saddlepoint approximation results (32), (40), (41), (50), (53), (57). The expressions for the corrections to the leading saddlepoint approximation require a different normalization factor. In our numerical experiments the restoration of the normalization did not lead to any systematic improvement.

In the case of the single portfolio, the representation of the saddlepoint approximation in terms of a shifted normal distribution does not really open any new possibilities because the saddlepoint integrals were computable exactly. The situation is different for multi-portfolio applications considered in the next section. We may not always be able to compute the saddlepoint integral, or even find the exact saddlepoint, but we may be able to improve the accuracy of the multivariate normal approximation by choosing it judiciously.

4.6 Saddlepoint and Monte Carlo importance sampling

We conclude the section by a brief discussion of the relationship between the saddlepoint analytics and the importance sampling techniques for direct Monte Carlo evaluation of credit risk measures discussed by Glasserman (2004). In the context of importance sampling with a single twisting parameter, the saddlepoint equation arises as an equation for the value of the measure twisting parameter ξ_0 that is optimal for the purpose of sampling the total loss in the vicinity of a given value K,

$$\sum_a \frac{w_a \mu_a \exp(\xi_0 w_a)}{1 - \mu_a + \mu_a \exp(\xi_0 w_a)} = K. \quad (67)$$

The quantities

$$\tilde{\mu}_a(\mathbf{X}) = \frac{\mu_a \exp(\xi_0 w_a)}{1 - \mu_a + \mu_a \exp(\xi_0 w_a)} \quad (68)$$

are the twisted conditional default probabilities to be used in the optimized Monte Carlo simulation of the individual risk factors Y_a. The correct conditional average is restored by means of the factor $\exp(\mathcal{K}(\xi_0|\mathbf{X}) - \xi_0 \sum w_a U_a)$. In this way the effectiveness of the somewhat abstract procedure of contour deformation gets an intuitive explanation in terms of measure adjustment to move the peak of the loss distribution to the region that needs to be sampled more accurately. In light of this link, the Edgeworth expansion of the CGF around $\xi = 0$ (which includes the normal proxy as the leading approximation) can be compared to a straightforward Monte Carlo simulation without importance sampling. Searching for a saddlepoint clearly appears better way to go in any analytical framework.

5 Analytical expansions II: CDO^2 and portfolio of single tranche CDOs

5.1 Laplace transform for CDO-Squared

As with the case of a single CDO, we proceed directly to the evaluation of the stop-loss integral

$$E[(\mathcal{L}-K)^+|\mathbf{X}] = \int_0^\infty \int_{\mathbf{c}-i\infty}^{\mathbf{c}+i\infty} \frac{\prod dL_p \prod d\xi_p}{(2\pi i)^M} (\mathcal{L}-K)^+ \exp\left(\mathcal{K}(\xi_1,\dots,\xi_M|\mathbf{X}) - \sum_p \xi_p L_p\right), \tag{69}$$

where CDO^2 portfolio loss $\mathcal{L}$ is given by Eq. (25), and the multi-portfolio CGF $\mathcal{K}(\xi_1,\dots,\xi_M|\mathbf{X})$ is given by Eq. (13). The integrals over the loss variables L_p can be decoupled at a price of the introduction of an additional inverse Laplace transform

$$(\mathcal{L}-K)^+ = \frac{1}{2\pi i}\int_{c_0-i\infty}^{c_0+i\infty} \frac{\exp(y(\mathcal{L}-K))}{y^2} dy. \tag{70}$$

After this substitution the integration over the loss variables reduces to the product of one-dimensional integrals

$$\int_0^\infty dL_p \exp[y((L_p-k_p)^+ - (L_p-K_p)^+) - \xi_p L^{(p)}], \tag{71}$$

which are easily evaluated. As a result we obtain

$$E[(\mathcal{L}-K)^+|\mathbf{X}] = \int_{c_0-i\infty}^{c_0+i\infty} \int_{\mathbf{c}-i\infty}^{\mathbf{c}+i\infty} \frac{dy \prod d\xi_p}{(2\pi i)^{M+1}} \frac{\exp(-yK)}{y^2} \exp(\mathcal{K}(\xi_1,\dots,\xi_M|\mathbf{X})) \prod_p \left(\frac{1}{\xi_p} + \frac{y\left[\exp(y(K_p-k_p)-\xi_p K_p) - \exp(-\xi_p k_p)\right]}{(y-\xi_p)\xi_p}\right). \tag{72}$$

An immense reduction in the dimensionality of the integration has been achieved from the initial N, the number of names, to $M+1$, the number of base CDO tranches plus one auxiliary variable introduced to decouple the losses. However it is not straightforward to take practical advantage of this reduction because the remaining integrals are over complex plane contours. The dimensionality is still too large for analytical integration, and the oscillatory nature of the integrand makes it impossible to implement Monte Carlo techniques directly.

The possibility of simplifying the integral (72) using a multidimensional generalization of the saddlepoint method remains open. In our numerical experiments we were able to locate the saddlepoint in the space of variables y and ξ_p but failed to obtain a robust computational algorithm. We turn to a more practical approach in the next subsection.

5.2 Semianalytical solution for the normal proxy method

The normal proxy method of Shelton (2004) consists in computing of the conditional average

$$E\left[(\mathcal{L}(L_1, L_2, \cdots, L_M) - K)^+ \;|\mathbf{X}\right] \tag{73}$$

where

$$\mathcal{L}(L_1, L_2, \cdots, L_M) = \sum_p \left((L_p-k_p)^+ - (L_p-K_p)^+\right), \tag{74}$$

assuming that the joint distribution of portfolio losses L_p conditional on $\mathbf{X}$ is multivariate normal with the average

$$\Lambda_p(\mathbf{X}) = \sum_a \mu_a(\mathbf{X})\, w_{a,p} \tag{75}$$

and covariance matrix

$$C_{p\,p'}(\mathbf{X}) = \sum_a \mu_a(\mathbf{X})\,(1-\mu_a(\mathbf{X}))\, w_{a,p} w_{a,p'}. \tag{76}$$

As with the single tranche case, the normal proxy method is recovered from the expansion of the CGF up to quadratic terms jointly in all variables $\{\xi_p\}$ around 0. The appearance of the multivariate normal distribution is not related to the Gaussian choice for the copula but is a consequence of the central limit theorem. We further omit explicit dependence on the central factor $\mathbf{X}$ and focus on the calculation of the conditional average with the understanding that the entire procedure is repeated with all values of $\mathbf{X}$ necessary to get the unconditional average.

A direct computation of the average (73) over a multivariate normal distribution still requires a Monte Carlo simulation. The simulation noise can be significantly reduced if we split from each loss variable L_p a part that is not correlated with other portfolios, similar to the factor formulation (8) of the copula model. To this end we decompose the portfolio loss as

$$L_p = \Lambda_p + D_p + \sigma_p W_p. \tag{77}$$

Here D_p are correlated normal variables that capture systematic portfolio risk. The variables W_p are independent identically distributed $N(0,1)$ variables uncorrelated with $\{D_p\}$, with the meaning of idiosyncratic risk factor for each individual portfolio. The amplitudes of the idiosyncratic terms σ_p are restricted by the requirement for the covariance matrix $E[D_p D_{p'}]$ to have non-negative eigenvalues. This requirement can be satisfied by taking all amplitudes equal to the smallest eigenvalue of the covariance matrix $C_{pp'}$.

The average of the variable D_p is zero due to the term Λ_p in the definition (77). The off-diagonal part of the covariance matrix $E[D_p D_{p'}]$ is $C_{p,p'}$ while the variances are reduced,

$$E[D_p^2] = C_{pp} - \sigma_p^2. \tag{78}$$

This reduces the simulation noise in the variables D_p in comparison with the initial loss variables L_p and makes it possible to perform the integration over D_p using just a few points $(D_1^{(i)}, D_2^{(i)}, \ldots D_M^{(i)})$ in the M-dimensional space, adjusted to match the first and second moments.

The remaining problem is to compute the decorrelated average over independent Gaussian $N(0,1)$ variables W_p

$$E\left[(\mathcal{L}(a_1+\sigma_1 W_1, \cdots, a_M + \sigma_M W_M) - K)^+\right] \tag{79}$$

with shifted centers $a_p = \Lambda_p + D_p^{(i)}$ dependent on the integration points $\{D_p^{(i)}\}$. Introducing an auxiliary variable y by the formula (70) the expectation over each W_p is decoupled,

$$\begin{aligned} & E\left[(\mathcal{L}(a_1+\sigma_1 W_1, \ldots, a_M + \sigma_M W_M) - K)^+\right] \\ = \quad & \frac{1}{2\pi i}\int_{C^+} \frac{dy}{y^2}\exp(-yK)\prod_p E[e^{y((a_p+\sigma_p W_p - k_p)^+ - (a_p+\sigma_p W_p - K_p)^+)}] \\ = \quad & \frac{1}{2\pi i}\int_{C^+} \frac{dy}{y^2}\exp\left(\sum_p \log(\chi_p(y)) - yK\right), \end{aligned} \tag{80}$$

with

$$\begin{aligned}\chi_p(y) &= \mathcal{N}(z_p^{(1)}) + e^{y\,(K_p-k_p)}\,\mathcal{N}(-z_p^{(2)}) \\ &+ e^{y\,(a_p-k_p)+\frac{1}{2}y^2\,\sigma_p^2}\,\Big(\mathcal{N}(z_p^{(2)}-y\sigma_p) - \mathcal{N}(z_p^{(1)}-y\sigma_p)\Big)\,, \end{aligned} \tag{81}$$

where constants $z_p^{(1)}$ and $z_p^{(2)}$ are $z_p^{(1)} = \frac{k_p - a_p}{\sigma_p}$ and $z_p^{(2)} = \frac{K_p - a_p}{\sigma_p}$. The residual one-dimensional integral over y can be taken using the saddlepoint technique described in Sect. 4.2 with the role of $\mathcal{K}(\xi)$ played by $\chi(y) = \sum_p \log(\chi_p(y))$.

5.3 Numerical results for CDO^2

Here we compare numerical results obtained using a direct Monte Carlo simulation of the copula model (**MC**), conditional multivariate normal approximation (**NP-MC**), and the semianalytical technique described in the previous section (**NP-AN**).

The underlying pool consists of 375 names with 5 year CDS spreads scattered around the average of 89 bp with standard deviation 71 bp. Based on this pool are 8 portfolios of 100 names with average overlap around 25%. Attachment points for the underlying tranches are 5.5% and 6.5%. The first set of results is for the values of the differential stop-loss $(\mathcal{L}(t)-k)^+ - (\mathcal{L}(t)-K)^+$ at two time horizons, $t = 1$ and $t = 4$, for the bounds $k = 12\%$ and $K = 28\%$, assuming recovery rate 40% and one of three value of Gaussian copula correlations, 10%, 40%, or 80%. The stop-loss is normalized to the tranche width and plotted as a function of CPU time in Figs. (1–6). The runs were done on a single CPU 1.8MHz Pentium M processor. The error bars for simulation based methods are obtained from the statistics over 36 independent runs.

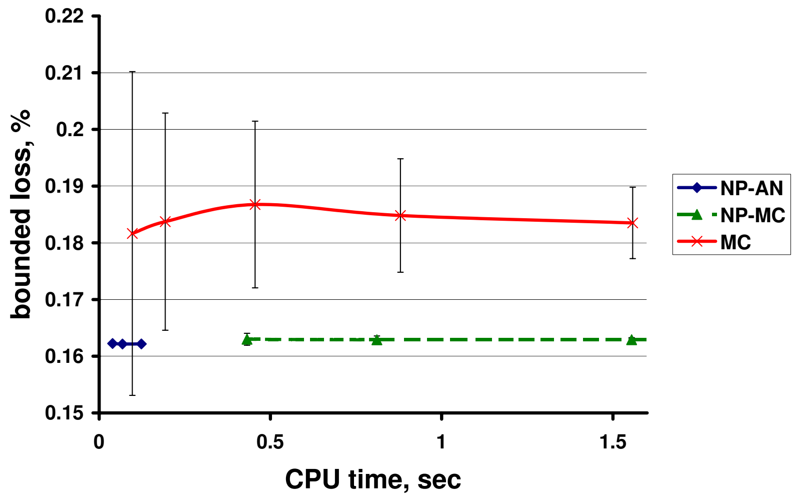

Figure 1: Normalized expected bounded loss for CDO^2: 1 year horizon and 10% correlation.

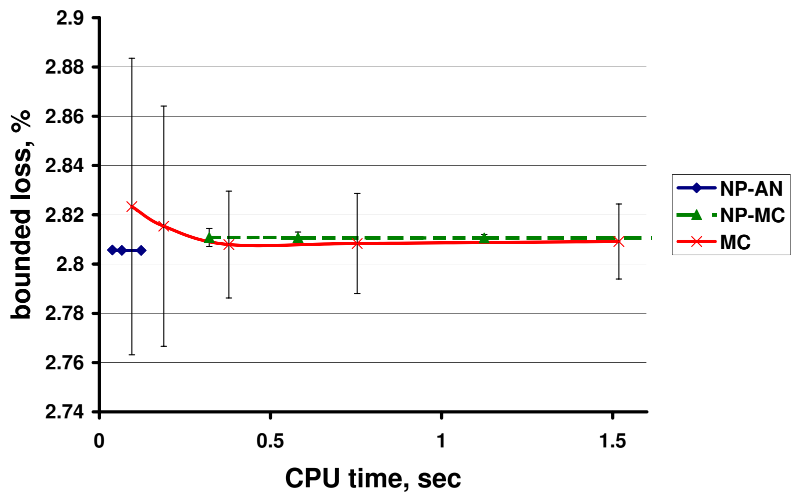

Figure 2: Normalized expected bounded loss for CDO^2: 1 year horizon and 40% correlation.

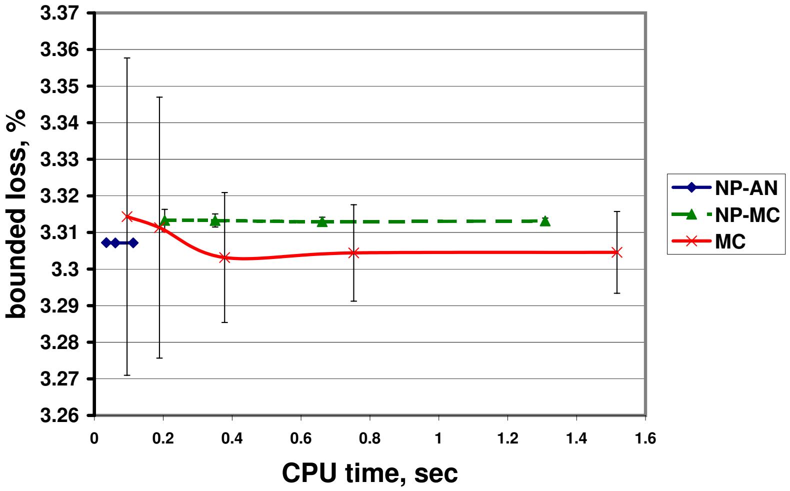

Figure 3: Normalized expected bounded loss for CDO^2: 1 year horizon and 80% correlation.

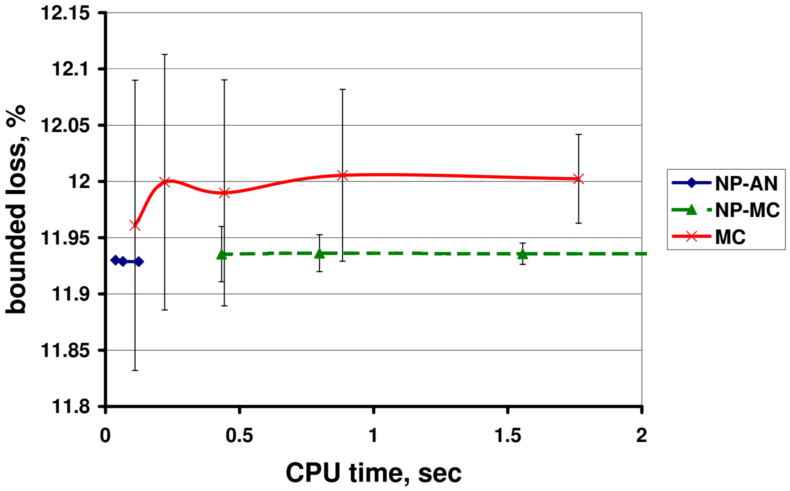

Figure 4: Normalized expected bounded loss for CDO^2: 4 year horizon and 10% correlation.

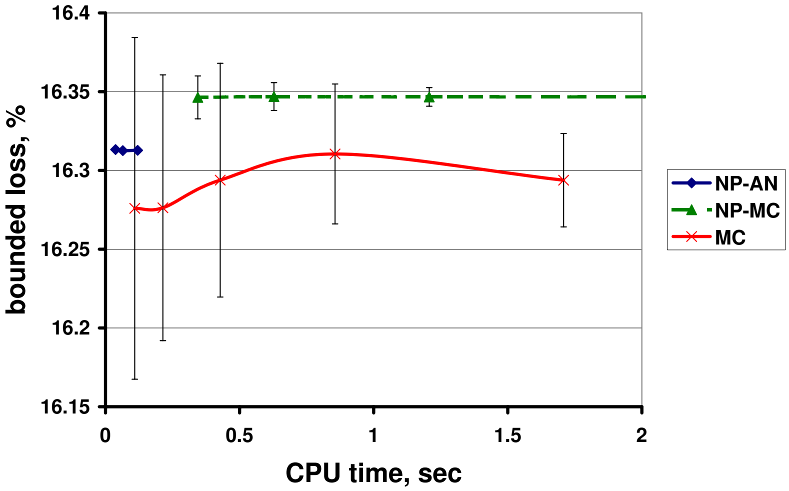

Figure 5: Normalized expected bounded loss for CDO^2: 4 year horizon and 40% correlation.

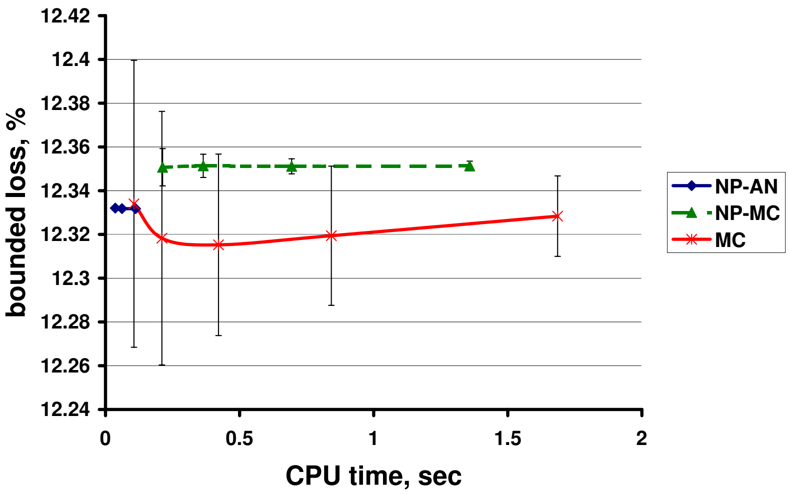

Figure 6: Normalized expected bounded loss for CDO^2: 4 year horizon and 80% correlation.

The next set of runs is for the quarterly par spread of a CDO^2 with 5 years to maturity and attachment points $k = 17.86\%$, $K = 28.57\%$, based on the same set of 8 mezzanine tranches.

CDO^2 par spread							
MC			**NP-MC**			**NP-AN**	
paths	value (bp)	time (sec)	paths	value (bp)	time (sec)	value (bp)	time (sec)
10000	406.02 ± 3.36	4.30	1000	407.82 ± 0.60	6.30	403.65	2.60
25000	405.59 ± 3.05	8.00	2500	407.89 ± 0.32	11.20	403.64	3.20
50000	405.72 ± 2.48	14.00	5000	407.89 ± 0.20	20.00	403.63	4.30

Finally we show results for hedge ratios to two arbitrarily chosen names out of the underlying pool. Name A has spread 91 bp and enters six portfolios, name B has spread 171 bp and enters one portfolio.

CDO^2 hedge ratio, name A							
MC			**NP-MC**			**NP-AN**	
paths	value (%)	time (sec)	paths	value (%)	time (sec)	value (%)	time (sec)
10000	5.93 ± 3.87	9.8	1000	5.61 ± 0.12	20.1	5.76	7.4
25000	5.26 ± 2.20	19.3	2500	5.64 ± 0.08	30	5.74	8.3
50000	5.23 ± 1.80	34	5000	5.62 ± 0.05	45	5.74	10.6

CDO^2 hedge ratio, name B							
MC			**NP-MC**			**NP-AN**	
paths	value (%)	time (sec)	paths	value (%)	time (sec)	value (%)	time (sec)
10000	0.14 ± 2.02	9.7	1000	0.80 ± 0.02	20	0.76	7.4
25000	0.52 ± 1.23	19.4	2500	0.81 ± 0.02	30.1	0.73	8.3
50000	0.59 ± 0.80	34.1	5000	0.81 ± 0.01	45.1	0.73	10.6

The results show that the bias introduced by the multivariate normal approximation is not worsened by further simplification described in the previous subsection. Thus the analytical approximation produces answers of the same accuracy in a fraction of time required to run a multivariate simulation.

5.4 Risk measures for portfolios of CDO tranches

Risk management for a portfolio of CDO tranches is based on details of the distribution of losses for the aggregate quantity $\mathcal{L}$, including the tail probability $P[\mathcal{L} \geq K]$ and expected shortfall $E(L|\mathcal{L} \geq K)$. Because of a simple relation between expected shortfall, tail probability, and stop-loss option, we only make a few remarks about the tail probability. The decoupling formula now takes the form

$$\theta(\mathcal{L} - K) = \frac{1}{2\pi i} \int_{c_0 - i\infty}^{c_0 + i\infty} \frac{\exp(y(\mathcal{L} - K))}{y} dy. \tag{82}$$

The only difference from the stop-loss expression is the power of y in the denominator,

$$P[\mathcal{L} \geq K|\mathbf{X}] = \int_{c_0-i\infty}^{c_0+i\infty} \int_{\mathbf{c}-i\infty}^{\mathbf{c}+i\infty} \frac{dy \prod d\xi_p}{(2\pi i)^{M+1}} \frac{\exp(-yK)}{y} \exp(\mathcal{K}(\xi_1, \ldots, \xi_M|\mathbf{X})) \prod_p \left(\frac{1}{\xi_p} + \frac{y\left[\exp(y(K_p - k_p) - \xi_p K_p) - \exp(-\xi_p k_p)\right]}{(y - \xi_p)\xi_p} \right). \quad (83)$$

The approximate techniques developed for the stop-loss apply with minor adjustments.

5.5 Failure of granularity adjustment for CDO^2

The derivation of the granularity adjustment for the stop-loss in the formula (63) carries over to the multidimensional case. This approach performed poorly in our numerical experiments. We only list here the final expression in the case of a one-dimensional central factor. With the notations introduced in Sect. 5.2,

$$\begin{aligned} E[(\mathcal{L} - K)^+] \;\approx\; & \int dx\, p(x)(\mathcal{L}(x) - K)^+ \\ & + \frac{1}{2} \sum_p \left[\frac{p(x_{pL}) C_{pp}(x_{pL}) I_L(x_{pL})}{|\Lambda_p'(x_{pL})|} - \frac{p(x_{pH}) C_{pp}(x_{pH}) I_L(x_{pH})}{|\Lambda_p'(x_{pH})|} \right] \\ & + \frac{1}{2} \sum_{p_1, p_2} \frac{p(x_0) C_{p_1 p_2}(x_0) \Theta_{k_{p_1}, K_{p_1}}(\Lambda_{p_1}(x_0)) \Theta_{k_{p_2}, K_{p_2}}(\Lambda_{p_2}(x_0))}{|\mathcal{L}'(x_0)|} \end{aligned} \quad (84)$$

where $\mathcal{L}(x) = \mathcal{L}(\Lambda_1(x), \ldots \Lambda_M(x))$, x_{pL} is the root of the equation $\Lambda_p(x) = k_p$, x_{pH} is the root of the equation $\Lambda_p(x) = K_p$, x_0 is the root of the equation $\mathcal{L}(x) = K$, $\Theta_{k_p, K_p}(x) = \theta(x - k_p) - \theta(x - K_p)$, and $I_L(x) = \theta(\mathcal{L}(x) - K)$.

6 Conclusions

We considered in a systematic way several analytical techniques for synthetic CDOs and credit default risk measures. Seemingly different methods of proxy normal integration, saddlepoint method, and granularity adjustment were recast in a unified form driven by the expansions of the cumulant generating function. Explicit expressions for the corrections to saddlepoint approximation for the tail probability, expected shortfall, and stop-loss option were put down in a form ready for numerical implementation. The decomposition of the VaR and expected shortfall into asset contribution in the leading saddlepoint approximation was also fully developed. We also made progress on the stop-loss problem for CDO^2, reducing the dimensionality of the integration and discussing conditional decorrelation of portfolios. The technique of decorrelation together with the saddlepoint evaluation of an integral over an auxiliary variable leads to results of the same accuracy as the multivariate normal approximation with significant runtime acceleration.

References

[1] L. Andersen, J. Sidenius, and S. Basu (2003) "All your hedges in one basket," Risk, November, 67–72.

[2] L. Andersen and J. Sidenius (2005) "Extensions to the Gaussian copula: random recovery and random factor loadings," Journal of Credit Risk, 1(1), 29–70.

[3] A. Berndt, R. Douglas, D. Duffie, M. Ferguson and D. Schranz (2005) "Measuring default risk premia from default swap rates and EDFs," BIS Working Paper No. 173.

[4] C. Bluhm, R. Overbeck, and C. Wagner (2003) "An introduction to credit risk modeling," Chapman & Hall/CRC, Chapter 6.

[5] H.-J. Brasch (2004) "A note on efficient pricing and risk calculation of credit basket products," working paper, available at www.defaultrisk.com.

[6] X. Burtschell, J. Gregory, and J.-P. Laurent (2005) "A comparative analysis of CDO pricing models," working paper, available at www.defaultrisk.com.

[7] H. E. Daniels (1954) "Saddlepoint approximations in statistics," The Annals of Mathematical Statistics, 25(4), 631–650.

[8] S. Emmer and D. Tasche (2005) "Calculating credit risk capital charges with the one-factor model," Journal of Risk, 7, pp. 85–101.

[9] P. Glasserman (2004) "Monte Carlo methods in financial engineering," Spinger-Verlag, Chapter 9.

[10] M. Gordy (2003) "A risk factor model foundation for ratings based capital rules," Journal of Financial Intermediation, 12(3), 199–232.

[11] C. Gourieroux, J.-P. Laurent, and O. Scaillet (2000) "Sensitivity analysis of values at risk," Journal of Empirical Finance, 7(3-4), 225–245.

[12] J. Hull, M. Predescu, and A. White (2005) "Bond prices, default probabilities, and risk premiums," Journal of Credit Risk, forthcoming.

[13] J. Hull and A. White (2004) "Valuation of a CDO and an n^{th} to default CDS without Monte Carlo simulation," Journal of Derivatives, 12(2), 8–23.

[14] R. Martin, K. Thompson, and C. Browne (2001a) "Unsystematic credit risk," Risk, June, 87–90.

[15] R. Martin, K. Thompson, and C. Browne (2001b) "Taking to the saddle," Risk, July, 91–94.

[16] R. Martin, K. Thompson, and C. Browne (2001c) "VaR: who contributes and how much," Risk, August, 99–102.

[17] R. Martin and T. Wilde (2002) "Unsystematic credit risk," Risk, November, 123–128.

[18] L. McGinty, E. Beinstein, R. Ahluwalia, and M. Watts (2004) "Introducing base correlations," JP Morgan, Credit Derivatives Strategy, London.

[19] D. X. Li (2000) "On default correlation: A copula function approach," The Journal of Fixed Income, 9(4), 43–54.

[20] R. Rebonato (2004) "Volatility and correlation," 2nd Ed., John Wiley & Sons Ltd, Chapter 1.

[21] P. Schönbucher (2002) "Taken to the limit: simple and not-so-simple loan loss distributions," working paper, avaiable at http://ssrn.com/abstract=378640.

[22] D. Shelton (2004) "Back to normal. Proxy integration: a fast accurate method for CDO and CDO-Squared pricing," Citigroup Global Structured Credit Research Report.

[23] D. Taras, C. Cloke-Browne, and E. Kalimtgis (2005) "Analytic improvement of the saddle-point approximation and spread risk attribution in a portfolio of tranches," Working paper, available at www.defaultrisk.com.

[24] T. Wilde (2001) "Probing granularity," Risk, August, 103–106.