

Quadratic Transform Approximation for CDO Pricing in Multifactor Models*

Paul Glasserman[†] and Sira Suchintabandit[‡]

Abstract. The multifactor version of copula models has the ability to generate complex correlation structure among defaults that is useful in fitting the base correlation skew. However, multifactor models have often been dismissed for their intractability. Even the semianalytical approach using Laplace transforms is computationally challenging, because although the model is tractable upon conditioning on the factors, unconditioning usually requires high efforts of integrating out the factors. To circumvent this problem, this paper develops a fast, closed-form approximation to the Laplace transform in multifactor models. The method, which approximates the conditional transform in a way that lends itself to closed-form unconditioning in arbitrarily high dimensions, is applicable to a range of models with Gaussian factors, including models that extend the standard Gaussian copula to allow stochastic recovery rates and factor loadings. We analyze the accuracy and convergence properties of the approximation. Numerical examples illustrate the speed and accuracy of the method.

Key words. CDOs, Laplace transforms, multifactor models, Gaussian copula, transform inversion, quadratic approximation

AMS subject classifications. 91G20, 91G40, 91G60

DOI. 10.1137/110827399

1. Introduction. The pricing of collateralized debt obligations (CDOs) and other portfolio credit derivatives requires the modeling of default times of obligors in the portfolio. One then seeks to compute the distribution of the portfolio's loss that results from defaults over fixed time horizons. Given these marginal loss distributions, valuing a CDO tranche reduces to calculating expectations of option-like payoffs tied to the portfolio losses. Calculating economic capital and other measures of portfolio credit risk also relies on a portfolio's loss distribution at fixed dates. Correlation among defaults plays an important role in determining the distribution of portfolio losses.

Correlation among reference entities is usually modeled by specifying a copula of default times. The copula is driven by a set of random “factors,” conditional on which the obligors become independent. The single-factor Gaussian copula has become a standard model for CDO pricing. In fact, the price of a CDO tranche is typically quoted through an *implied correlation* parameter associated with the single-factor Gaussian copula model. (Another popular measure of correlation is known as the *base correlation*, which is also implied from the single-factor Gaussian copula model.)

The single-factor Gaussian copula model, however, is not consistent with the market in that it cannot match the prices of all CDO tranches simultaneously. This is evident in the smile

*Received by the editors March 14, 2011; accepted for publication (in revised form) November 29, 2011; published electronically January 26, 2012.

<http://www.siam.org/journals/sifin/3/82739.html>

[†]Department of Decision, Risk, and Operations, Columbia Business School, New York, NY 10027 (pg20@columbia.edu).

[‡]Department of Banking and Finance, Chulalongkorn University, 10330 Bangkok, Thailand (sira@acc.chula.ac.th).

of the implied correlation curve and the skew of the base correlation curve. Various models have been proposed in the attempt to fit this correlation skew. Some replace the Gaussian factor with more heavy-tailed distribution (see, for example, Burtshell, Gregory, and Laurent [6] and Gregory and Laurent [11]), while others extend the Gaussian copula model to allow randomness in the recovery rates and factor loadings (see, for example, Andersen and Sidenius [5] and Burtshell, Gregory, and Laurent [7]). However, these models are mostly tested in their single-factor versions, while their multifactor counterparts have not received much treatment, most likely because of the added complexity in computing the model outputs. In particular, the multifactor Gaussian copula model, the most natural extension to the standard single-factor model, has been left largely unexplored.

Multifactor models are appealing for several reasons. First, they enable one to create a richer correlation structure that offers more flexibility in calibrating to the market correlation skew. As we find in our numerical examples, it is possible to generate the skew simply by including some “group-specific factors”—factors that affect only certain subsets of obligors in the portfolio. Multifactor models with group-specific factors can also, in theory, find use in problems that involve more than one CDO structure; in such problems, the correlation ascribed to a pair of names can vary depending on the CDO structure to which they belong. Last, there is a conceptual link between multifactor models and Merton’s structural model of default [19]; the latent variables in the normal copula can be associated with changes in the values of the underlying firms. In such an idealized setting, the number of factors describing the correlations between these variables would be comparable to the number of factors in an equity valuation model, which could easily be 10 or more.

Pricing under multifactor models, however, is challenging. The standard procedure typically begins by conditioning on the factors, after which obligors become independent. This enables one to compute the conditional loss distribution through convolution of the individual loss distributions (as in the recursive method of Andersen, Basu, and Sidenius [4]). Alternatively, one computes the Laplace transform of the conditional loss distribution, which, thanks to the conditional independence, is simply the product of the transforms of losses from individual assets (see, for example, Gregory and Laurent [10], Haaf, Reiss, and Schoenmakers [12], and Laurent and Gregory [15]). (Yet another approach, presented in Iscoe et al. [14], uses a representation of the hockey stick function to directly approximate tranche prices, conditional on the factors.) All the approaches mentioned above ultimately require one to uncondition the factors, which entails multidimensional integration in as many dimensions as there are factors. This quickly becomes infeasible when the number of factors is high, and it presents the main obstacle in the practical use of multifactor models.

This paper develops a *quadratic transform approximation (QTA)* technique to address this problem. Our method approximates the conditional Laplace transform—which is readily available—with the exponential of a quadratic function of the factors. We choose this class of functions because their expectations can be evaluated in closed form, regardless of the number of factors. We thus circumvent the difficulty of the high-dimensional integration required to find the unconditional transform as follows: rather than carry out an approximate integration of the exact conditional transform (as one would in the traditional approaches of numerical integration or Monte Carlo simulation), we do an exact integration of an approximation to the conditional transform. This method results in a fast, closed-form approximation to the

unconditional Laplace transform, which can then be used in a Laplace inversion procedure to obtain the loss distribution in multifactor models.

We provide a rigorous analysis of the convergence property of the method. We bound the error in our approximation in terms of a norm $\|\mathbf{A}\|$ of the factor loading matrix \mathbf{A} in the multifactor Gaussian copula framework. The loading matrix determines the strength of the correlations between names in the portfolio; the off-diagonal entries of $\mathbf{A}\mathbf{A}^\top$ are the correlations between distinct pairs of names in the portfolios. We show that our method calculates tranche prices and loss probabilities with an error that is $O(\|\mathbf{A}\|^4)$; remarkably, this can be improved to $O(\|\mathbf{A}\|^6)$ with a simple, strategic choice of parameters that eliminates the leading order error term.

As these results suggest, the method is most accurate when the loading coefficients (and thus the correlations) are not too large, but it becomes coarse when correlations are high. However, the method can be easily refined to handle the realistic case of strong correlation. Typically, stronger correlations arise primarily through one of the factors, which is then interpreted as a “market” factor. For correlations of this form, we can either condition on the market factor or segment it and integrate each segment separately. These extra steps require computing time that is similar to that in the single-factor case; in other words, our procedure offers a way to handle multifactor models at a computational cost comparable to single-factor models. Our numerical examples show that the method is quite accurate, even in the environment of strong correlation.

It should be noted that the QTA method is applicable to a range of models that have Gaussian factors, which include models that extend the standard Gaussian copula to allow randomness in the recovery rates and factor loadings (see Andersen and Sidenius [5]). Although the single-factor versions of these models are already known to have desirable traits such as fat tails and skew-generating property, their multifactor counterparts, as shown in our numerical examples, offer more flexibility when one attempts to fit real market data that exhibit strong correlation and pronounced skew. The numerical examples also demonstrate the accuracy of the QTA method as it is applied to this class of models.

The organization of this paper is as follows. In section 2, we review the normal copula model and the method of Laplace inversion in risk measurement and derivative pricing. Section 3 is devoted to describing and analyzing our approximation method. Section 4 explains the extension to strong correlations among obligors and provides numerical examples. Proofs of our main results are deferred to the appendices.

2. Background.

2.1. Normal copula model for credit risk. The valuation of CDOs relies on the modeling of the reference portfolio’s *loss process*, whose value at each time represents the accumulation of default losses of credit instruments in the portfolio. This section provides a description of, and explains the model incentive for, the multifactor version of the normal copula model (Li [16]). Extensions to more elaborated models in the Gaussian family will be discussed in section 5.

Suppose that the reference portfolio comprises N obligors. For $j = 1, \dots, N$, let τ_j denote the time that the j th obligor defaults, and let Y_j denote loss that results from the default of

the j th obligor. Let $L(t)$ denote the aggregated default losses at time t . Thus

$$(2.1) \quad L(t) = \sum_{j=1}^N Y_j 1\{\tau_j \leq t\}.$$

We assume that τ_j is a random variable whose marginal distribution function, $F_j(t) = \mathbb{P}(\tau_j \leq t)$, is known (e.g., from the quotes of credit default swap spreads at different maturities). The loss-given-default Y_j represents the part of the notional principle of obligor j that cannot be recovered. To reflect uncertainty in the recovery rate, Y_j is allowed to be random, with the distribution that may depend on time. Let us assume for the time being that Y_j 's are independent both of each other and of the default times τ_j . (We will relax this assumption in section 5.)

To sample a correlated set of default times τ_1, \dots, τ_N with marginal distributions F_1, \dots, F_N , the normal copula model specifies the default times as

$$(2.2) \quad \tau_j = F_j^{-1}(\Phi(X_j)), \quad j = 1, 2, \dots, N,$$

where Φ denotes the cumulative normal distribution function, and X_j ($j = 1, \dots, N$) are *correlated* $\mathcal{N}(0, 1)$ random variables. In the multifactor Gaussian copula model, correlation among X_j is introduced through a $d \times 1$ vector $\mathbf{Z} = [Z_1 \ Z_2 \ \dots \ Z_d]^\top$ of independent standard normal random variables:

$$(2.3) \quad X_j = \mathbf{a}_j^\top \mathbf{Z} + b_j \epsilon_j, \quad i = 1, \dots, N.$$

Here, $\mathbf{a}_j = [a_{j1} \ a_{j2} \ \dots \ a_{jd}]^\top$ is a $d \times 1$ vector of real constants satisfying $\mathbf{a}_j^\top \mathbf{a}_j < 1$, and ϵ_j are $\mathcal{N}(0, 1)$ random variables, independent of each other and independent of \mathbf{Z} . The constant b_j is chosen so that X_j has unit variance.

The incentive for considering the multifactor version of the normal copula model becomes clear when one rewrites (2.3) in matrix form:

$$(2.4) \quad \begin{bmatrix} X_1 \\ X_2 \\ \vdots \\ X_N \end{bmatrix} = \begin{bmatrix} a_{11} \\ a_{21} \\ \vdots \\ a_{N1} \end{bmatrix} Z_1 + \begin{bmatrix} a_{12} \\ a_{22} \\ \vdots \\ a_{N2} \end{bmatrix} Z_2 + \dots + \begin{bmatrix} a_{1d} \\ a_{2d} \\ \vdots \\ a_{Nd} \end{bmatrix} Z_d + \begin{bmatrix} b_1 \epsilon_1 \\ b_2 \epsilon_2 \\ \vdots \\ b_M \epsilon_N \end{bmatrix}.$$

While each ϵ_j represents the idiosyncratic factor affecting only the j th obligor, each Z_j ($j = 1, \dots, d$) can affect all (or a certain group of) obligors. Although the factors Z_j are sometimes given economic interpretations (as industry or regional risk factors, for example), the key role of the factors Z_j is that they allow us to model complicated correlation structure in a nonhomogeneous portfolio.

For example, suppose that obligors are classified into two groups: the high-loss group $\mathcal{D} \subset \{1, \dots, N\}$, consisting of obligors with low recovery rates, and the low-loss group \mathcal{D}^c , consisting of obligors with high recovery rates. Suppose that, for the high-loss group \mathcal{D} , we

designate Z_k as the \mathcal{D} -specific factor, with its loading given by $a_{jk} = \sqrt{\varrho}$ for all $j \in \mathcal{D}$ and $a_{jk} = 0$ for all $j \notin \mathcal{D}$. The effect of the intragroup correlation ϱ on the prices of CDO tranches is highly nonlinear and varies across different tranches. This is advantageous in producing correlation skew observed in the market (as one will see in the numerical examples in sections 3.3, 4.3, and 5.3). In richer models with more than two groups, intra- and intergroup correlations can be modeled simply by including more factors.

The matrix $\mathbf{A} := [a_{ij}]_{N \times d}$ is referred to as the *loading matrix*. For the time being, we assume that \mathbf{A} is constant. (The following analysis holds even when \mathbf{A} is a deterministic function of time.) In section 5, we will discuss the case when \mathbf{A} is stochastic.

2.2. Pricing through Laplace transforms. Given the model for $L(t)$, the pricing problem reduces to computing the distribution of $L(t)$ at a set of fixed times t , as we now explain. (The following paragraph contains an argument that is explained in more detail in Hull and White [13].)

The cash flow of a CDO can be decomposed as follows. *Coupon payments* (seller to buyer) are made periodically based on the remaining notional principal at coupon dates. *Default payments* (buyer to seller), which reduce the notional principal, occur whenever reference entities default. *Accrual payments* (seller to buyer) are based on the reduction in the notional principal that takes place in between coupon dates. Suppose that we discretize time as t_0, t_1, t_2, \dots . Consider a CDO tranche whose lower and upper attachment points are A and B . The expected loss absorbed by this tranche during the period $(t_k, t_{k+1}]$ is

$$(2.5) \quad \left(\mathbb{E}[B \wedge L(t_{k+1})] - \mathbb{E}[A \wedge L(t_{k+1})] \right) - \left(\mathbb{E}[B \wedge L(t_k)] - \mathbb{E}[A \wedge L(t_k)] \right).$$

If we can compute this quantity for every period $(t_k, t_{k+1}]$ during the life of the CDO, then the expectation of all cash flows will be determined. Therefore, the problem of valuing a CDO reduces to the problem of computing expectation of the form $\mathbb{E}[y \wedge L(t)]$ at fixed points t in time. So, from now on, we will consider the loss at fixed t and simply write $L(t)$ as L .

Let $i = \sqrt{-1}$ denote the imaginary unit, so that every complex number s can be decomposed as $s = \operatorname{Re} s + i \operatorname{Im} s$. Let \mathbb{C}^+ denote the set of all complex number s with $\operatorname{Re} s \geq 0$. The Laplace transform of the portfolio loss is the mapping

$$\phi(s) \quad := \quad \mathbb{E}[e^{-sL}], \quad s \in \mathbb{C}^+.$$

The distribution of L can be obtained from $\phi(s)$ through the process of Laplace inversion (for background, see, for example, Abate and Whitt [2]). In particular, the expectation of the form (2.5) can be obtained from either one of the following inversion integrals:

$$(2.6) \quad \mathbb{E}(L \wedge y) = \frac{1}{i\pi} \int_{a-i\infty}^{a+i\infty} \operatorname{Re}(e^{sy}) \operatorname{Re}\left(\frac{1 - \phi(s)}{s^2}\right) ds$$

$$(2.7) \quad = \frac{1}{i\pi} \int_{a-i\infty}^{a+i\infty} \operatorname{Re}(e^{sy} - e^{ay}) \operatorname{Re}\left(\frac{1 - \phi(s)}{s^2}\right) ds,$$

where integration is on the contour $\operatorname{Re} s = a > 0$ in the complex plane. While these two integrals can be easily shown to be equivalent, the integral (2.7) has an advantage in that,

when $a = 0$, its integrand is more continuous around the point $s = 0$. (For this reason, we will use (2.7) for Laplace inversion in all our subsequent numerical examples.) Note that similar forms of inversion integrals arise frequently in option pricing; see, for example, Carr and Madan [8] and Lord and Kahl [17]. Such integrals are also used for saddlepoint approximations, as in Martin, Thompson, and Browne [18] and Gordy [9].

In practice, Laplace inversion is carried out numerically. For example, Abate and Whitt [2] apply a trapezoidal rule with stepsize $h = \pi/2y$ to the integral (2.6), using $a = A/y$ as the integrating contour (A is a positive real number that is used to control the discretization error). The result is a fast-converging, nearly alternating series:

$$\mathbb{E}(L \wedge y) \approx \frac{e^A y}{2} \sum_{k=-\infty}^{\infty} (-1)^k \operatorname{Re} \frac{1 - \phi(A/y + i\pi k/y)}{(A + i\pi k)^2}.$$

Other examples of inversion methods involve the *characteristic function*, which is defined as the mapping $\omega \mapsto \phi(-i\omega)$, where $\omega \in \mathbb{R}$. We refer the reader to Abate and Whitt [1] for background on inverting characteristic functions.

Numerical inversion techniques reduce the problem of computing the distribution of the tranche losses to the problem of evaluating the Laplace transform $\phi(s)$. We now explain the typical procedure for calculating the Laplace transform.

Using the factor structure (2.3), obligors become independent conditional on \mathbf{Z} . Using (2.2) and (2.3) to invert the relationship between τ_j and the standard Gaussian random variable ϵ_j , one obtains the conditional probability of default within some fixed time t :

$$(2.8) \quad \mathbb{P}(\tau_j \leq t | \mathbf{Z}) = \Phi\left(\frac{x_j - \mathbf{a}_j^\top \mathbf{Z}}{b_j}\right), \quad \text{where } x_j = \Phi^{-1}(F_j(t)).$$

Let $\psi_j(s) = \mathbb{E}[e^{-sY_j}]$ be the Laplace transform of Y_j , which is assumed to be known explicitly for all $s \in \mathbb{C}^+$. Using conditional independence,

$$(2.9) \quad \begin{aligned} \mathbb{E}[e^{-sL} | \mathbf{Z}] &= \prod_{j=1}^N \mathbb{E}\left[e^{-sY_j} 1_{\{\tau_j \leq t\}} \mid \mathbf{Z}\right] \\ &= \prod_{j=1}^N \left(\mathbb{P}(\tau_j > t | \mathbf{Z}) \mathbb{E}[e^{-s \cdot 0}] + \mathbb{P}(\tau_j \leq t | \mathbf{Z}) \mathbb{E}[e^{-sY_j}] \right) \\ &= \prod_{j=1}^N \left(1 + (\psi_j(s) - 1) \Phi\left(\frac{x_j - \mathbf{a}_j^\top \mathbf{Z}}{\sqrt{1 - \mathbf{a}_j^\top \mathbf{a}_j}}\right) \right). \end{aligned}$$

(Here, we assume that Y_j is independent of τ_j . We will relax this assumption in section 5.)

This conditional Laplace transform depends on time through the definition of x_j (see (2.8)). Dependence on time may also enter through ψ_j and \mathbf{a}_j if the loss Y_j and the loading matrix \mathbf{A} are specified as time-varying. However, we choose not to make the dependence on time explicit in our notation, for the sake of simplicity.

From (2.9), the Laplace transform $\phi(s) = \mathbb{E}[e^{-sL}]$ is obtained by integrating over the factors \mathbf{Z} . In principle, unconditioning can be achieved by, say, numerical integration or

quasi-Monte Carlo simulation. These methods, however, are inappropriate for the multifactor setting (2.3), because computing time grows exponentially with the number of factors. This prompts the need for a faster method for computing $\phi(s)$, particularly seeing as Laplace inversion requires multiple evaluation of $\phi(s)$ at many values of s .

Our main contribution, and the focus of the rest of the paper, is the development of an efficient, analytical method for approximating $\phi(s)$. The key idea in our approach is to approximate the conditional transform (2.9) in a way that facilitates integration in arbitrarily high dimensions. To put it another way, rather than attempt an approximate integration of the exact conditional transform, we carry out an exact integration of an approximation to the conditional transform.

3. Approximating the Laplace transform. In this section, we propose and analyze our quadratic transform approximation (QTA) for the Laplace transform $\phi(s)$. The description of the method is given in section 3.1 (for a preview, see Steps 1–3 at the end of section 3.1). In section 3.2, we state some convergence properties of the approximation.

3.1. Description of the approximation. We begin by stating a proposition that underlies our approximation. The relevance of the proposition is that it identifies a class of functions of normal random vectors whose expectations can be evaluated in closed form, regardless of the dimension of the problem.

Proposition 1. *Let \mathbf{Z} be a $d \times 1$ vector of independent standard normal variables. For any scalar $c \in \mathbb{C}$, vector $\mathbf{g} \in \mathbb{C}^d$, and matrix $\mathbf{H} \in \mathbb{C}^{d \times d}$ for which $\text{Re } \mathbf{H}$ is negative-semidefinite,*

$$(3.1) \quad \mathbb{E} \left[e^{c + \mathbf{g}^\top \mathbf{Z} + \mathbf{Z}^\top \mathbf{H} \mathbf{Z}} \right] = \frac{1}{\sqrt{\det(\mathbf{I} - 2\mathbf{H})}} e^{c + \mathbf{g}^\top (\mathbf{I} - 2\mathbf{H})^{-1} \mathbf{g}/2}.$$

This proposition forms the basis for our approach to approximate the Laplace transform $\phi(s) = \mathbb{E}[e^{-sL}]$. The main idea is to approximate the conditional Laplace transform (2.9) by an exponential of a quadratic function of \mathbf{Z} and then use Proposition 1 to uncondition the risk factors.

In describing our approximation scheme, the following notations are useful. For a fixed $s \in \mathbb{C}$, define the mappings $v \mapsto g_j(v)$ ($j = 1, \dots, N$) as

$$(3.2) \quad g_j(v) := 1 + (\psi_j(s) - 1) \Phi \left(\frac{x_j + v \sqrt{\mathbf{a}_j^\top \mathbf{a}_j}}{b_j} \right), \quad v \in \mathbb{R},$$

where $\psi_j(s)$ is the Laplace transform of the loss-given-default Y_j . With this notation, we rewrite (2.9) as

$$(3.3) \quad \mathbb{E} [e^{-sL} \mid \mathbf{Z}] = \prod_{j=1}^N g_j(V_j) = e^{\sum_{j=1}^N \ln g_j(V_j)}, \quad \text{where } V_j := -\frac{\mathbf{a}_j^\top \mathbf{Z}}{\sqrt{\mathbf{a}_j^\top \mathbf{a}_j}}.$$

(In this paper, $\ln x$, where $x \in \mathbb{C}$, denotes the unique complex number that satisfies $\exp(\ln x) = x$ and $-\frac{\pi}{2} \leq \text{Im}(\ln x) \leq \frac{\pi}{2}$.)

The Laplace transform $\phi(s) = \mathbb{E}[e^{-sL}]$ is obtained by unconditioning (3.3), which usually requires multidimensional integration with respect to \mathbf{Z} . If, however, we can approximate

the exponent $\sum_{j=1}^N \ln g_j(V_j)$ by a quadratic function of \mathbf{Z} , then we can use Proposition 1 to uncondition (3.3) and thus obtain a closed-form approximant to $\phi(s)$. We propose that each $\ln g_j(V_j)$ in the exponent be approximated by

$$(3.4) \quad \ln g_j(V_j) \approx \alpha_j + \beta_j V_j + \eta_j V_j^2,$$

where the scalars α_j , β_j , and η_j are complex-valued. Using this approximation,

$$(3.5) \quad \phi(s) = \mathbb{E} \left[e^{\sum_{j=1}^N \ln g_j(V_j)} \right] \approx \mathbb{E} \left[e^{\sum_{j=1}^N (\alpha_j + \beta_j V_j + \eta_j V_j^2)} \right] = \mathbb{E} \left[e^{c + \mathbf{g}^\top \mathbf{Z} + \mathbf{Z}^\top \mathbf{H} \mathbf{Z}} \right],$$

where the last equality follows from the fact that V_j 's are linear in \mathbf{Z} (see (3.3)). The scalar b , the vector \mathbf{g} , and the matrix \mathbf{H} are given explicitly by

$$(3.6) \quad c = \sum_{j=1}^N \alpha_j, \quad \mathbf{g} = - \sum_{j=1}^N \frac{\beta_j \mathbf{a}_j}{\sqrt{\mathbf{a}_j^\top \mathbf{a}_j}}, \quad \text{and} \quad \mathbf{H} = \sum_{j=1}^N \eta_j \frac{\mathbf{a}_j \mathbf{a}_j^\top}{\mathbf{a}_j^\top \mathbf{a}_j}.$$

The last expectation in (3.5) provides an approximant for the Laplace transform $\phi(s)$. Proposition 1 is used to evaluate the expectation in closed form.

To complete the description of our method, it is left only to explain how to obtain the coefficients $(\alpha_j, \beta_j, \eta_j)$ in the approximation (3.4). We use the weighted least-squares method to fit the quadratic function (3.4); that is to say, we choose α_j , β_j , and η_j that solve the minimization problem

$$(3.7) \quad \min_{\lambda \in \Lambda} \sum_{\lambda \in \Lambda} f(\lambda) \left| \ln g_j(\lambda) - \alpha_j - \beta_j \lambda - \eta_j \lambda^2 \right|^2.$$

The summands represent the approximation errors at certain gridpoints $\lambda \in \Lambda$, where $f(\lambda)$ represents the penalty weight for the errors. We assume that Λ and f are the same for all j and that $\sum_{\lambda \in \Lambda} f(\lambda) = 1$. As the rationale behind (3.7) is to minimize the *expected* error of the approximation (3.4) over possible realizations of V_j , the set of gridpoints Λ and the weight f should be chosen to reflect the fact that V_j is standard Gaussian. For our numerical example, we arrange the gridpoints evenly between -3 and 3 , and specify $f(\lambda)$ as exponentially decreasing in $|\lambda|^2$, in accordance with the normal distribution. While the alternative of using the unweighted (i.e., equally weighted) least-squares scheme also produces acceptable approximation $\hat{\phi}(s)$ in our numerical experiment, we find that the weighted scheme yields a noticeably better result. Indeed, the convergence result in section 3.2 confirms the appropriateness of specifying the weighting function $f(\lambda)$ to comply with the normal density function (in the sense that the first six moments are matched; see Theorem 1).

The advantage of using the least-squares method to determine α_j , β_j , and η_j is that the optimization problem (3.7) has a unique, closed-form solution. To characterize the solution, define $\bar{\lambda}^n := \sum_{\lambda \in \Lambda} \lambda^n f(\lambda)$. Then, the solution to (3.7) is given by

$$(3.8) \quad \begin{bmatrix} \alpha_j \\ \beta_j \\ \eta_j \end{bmatrix} = \begin{bmatrix} 1 & \bar{\lambda} & \bar{\lambda}^2 \\ \bar{\lambda} & \bar{\lambda}^2 & \bar{\lambda}^3 \\ \bar{\lambda}^2 & \bar{\lambda}^3 & \bar{\lambda}^4 \end{bmatrix}^{-1} \sum_{\lambda \in \Lambda} \begin{bmatrix} f(\lambda) \ln g_j(\lambda) \\ \lambda f(\lambda) \ln g_j(\lambda) \\ \lambda^2 f(\lambda) \ln g_j(\lambda) \end{bmatrix}.$$

At this point we have given the full description of the method of approximating the Laplace transform $\phi(s)$. We now summarize the steps in the procedure.

Summary. *The QTA method to approximate the Laplace transform $\phi(s)$ at a fixed value of s is as follows:*

- Step 1 For each $j = 1, \dots, N$, compute α_j , β_j , and η_j from (3.8).
- Step 2 Compute c , \mathbf{g} , and \mathbf{H} from (3.6).
- Step 3 The approximant for $\phi(s)$ is

$$\widehat{\phi}(s) := \mathbb{E} \left[e^{c + \mathbf{g}^\top \mathbf{Z} + \mathbf{Z}^\top \mathbf{H} \mathbf{Z}} \right] = \frac{1}{\sqrt{\det(\mathbf{I} - 2\mathbf{H})}} e^{c + \mathbf{g}^\top (\mathbf{I} - 2\mathbf{H})^{-1} \mathbf{g}/2}.$$

This procedure provides a fast, analytical way of approximating the Laplace transform $\phi(s)$, since all steps are in closed form and no numerical integration is required. The approximant $\widehat{\phi}(s)$ can then be used in place of $\phi(s)$ in the inversion integrals (from section 2.2) to compute an approximation for the tranche price $\mathbb{E}(L - y)^+$ or other quantities of interest.

To end this section, we note that the validity of the formula (3.5) is guaranteed when $\text{Re } \mathbf{H}$ is negative-semidefinite. A convenient way to ensure negative-semidefiniteness of $\text{Re } \mathbf{H}$ is to impose that $\text{Re } \eta_j \leq 0$ for all j (see (3.6)). If, for some j , Step 1 returns η_j with $\text{Re } \eta_j > 0$, then we reset $\text{Re } \eta_j$ to zero and refit the real part of $\ln g_j(V_j)$ in (3.4) by a linear (instead of quadratic) function of V_j . When applied to the market data, however, Step 1 hardly ever returns η_j that violates $\text{Re } \eta_j \leq 0$. For example, if the marginal default probability of name j is less than 20%, then $x_j = \Phi^{-1}(0.2) > 0.84$ and $\text{Re}(\ln g_j(v))$ can be shown to be concave for all $v > -(x_j - 0.84)/\sqrt{\mathbf{a}_j^\top \mathbf{a}_j}$. This implies that η_j in (3.4) satisfies $\text{Re } \eta_j \leq 0$.

3.2. Convergence theorem. The accuracy of $\widehat{\phi}(s)$ depends on the goodness of the approximation (3.4). The goodness of fit depends primarily on the magnitude of the loadings in \mathbf{A} . If $\sqrt{\mathbf{a}_j^\top \mathbf{a}_j}$ is small, then it can be seen from (3.2) that $\ln g_j(V_j)$ will be almost linear in V_j , and therefore (3.4) will fit better. (In the extreme case where $\sqrt{\mathbf{a}_j^\top \mathbf{a}_j} = 0$, $\ln g_j(V_j)$ becomes constant for all V_j and the approximation (3.4) becomes exact.) Therefore, it can be said that the approximation error $|\phi(s) - \widehat{\phi}(s)|$ decreases along with the norm of \mathbf{A} . In fact, if we let $\|\mathbf{A}\|$ denote the ∞ -norm of \mathbf{A} (defined as $\|\mathbf{A}\| = \max_j \sum_k |a_{jk}|$), then it is easy to show that $|\phi(s) - \widehat{\phi}(s)| \rightarrow 0$ as $\|\mathbf{A}\| \rightarrow 0$. In other words, the approximation $\widehat{\phi}(s)$ becomes exact when $\|\mathbf{A}\| \rightarrow 0$.

This section poses the following question: How fast does $\widehat{\phi}(s)$ converge to $\phi(s)$? Theorem 1 answers this question by giving the rate of convergence. The proof is given in Appendix B.

Theorem 1. *Assume that the marginal default probabilities of obligors are less than 1/2:*

(i) *Suppose that the gridpoints $\lambda \in \Lambda$ are arranged symmetrically around zero (so that $\overline{\lambda^m} = 0$ whenever m is odd). Then, there exist real constants δ and C such that for all loading matrices \mathbf{A} with $\|\mathbf{A}\| < \delta$,*

$$(3.9) \quad \left| \phi(s) - \widehat{\phi}(s) \right| < C \|\mathbf{A}\|^4 \quad \forall s \in \mathbb{C}^+.$$

In other words, $\widehat{\phi}(s)$ converges to $\phi(s)$ at the rate of $\|\mathbf{A}\|^4$.

(ii) *Suppose, in addition, that $\overline{\lambda^2}$, $\overline{\lambda^4}$, and $\overline{\lambda^6}$ match the second, fourth, and sixth moments, respectively, of the standard normal random distribution; that is to say, $\overline{\lambda^2} = 1$, $\overline{\lambda^4} = 3$, and*

$\bar{\lambda}^6 = 15$. Then, there exist real constants δ' and C' such that for all loading matrices \mathbf{A} with $\|\mathbf{A}\| < \delta'$,

$$(3.10) \quad \left| \phi(s) - \widehat{\phi}(s) \right| < C' \|\mathbf{A}\|^6 \quad \forall s \in \mathbb{C}^+.$$

In other words, $\widehat{\phi}(s)$ converges to $\phi(s)$ at the rate of $\|\mathbf{A}\|^6$. ■

Next, we will show that the convergence rates in Theorem 1 are preserved in the process of Laplace transform inversion. Let $F(y)$ be the approximant of $\mathbb{E}(L \wedge y)$ obtained by replacing $\phi(s)$ in (2.7) with $\widehat{\phi}(s)$; that is,

$$F(y) := \frac{1}{i\pi} \int_{a-i\infty}^{a+i\infty} \operatorname{Re}(e^{sy} - e^{ay}) \operatorname{Re}\left(\frac{1 - \widehat{\phi}(s)}{s^2}\right) ds.$$

Corollary 1. *Under the assumptions of Theorem 1, $F(y)$ converges to $\mathbb{E}(L-y)^+$ at the rate of $\|\mathbf{A}\|^4$ for all real y . If, in addition, $\bar{\lambda}^2 = 1$, $\lambda^4 = 3$, and $\bar{\lambda}^6 = 15$, then $F(y)$ converges to $\mathbb{E}(L-y)^+$ at the rate of $\|\mathbf{A}\|^6$ for all real y .*

Proof. Assume the premise of Theorem 1. From the definition of $F(y)$ and from (2.7),

$$\begin{aligned} |F(y) - \mathbb{E}(L \wedge y)| &= \left| \frac{1}{i\pi} \int_{a-i\infty}^{a+i\infty} \operatorname{Re}(e^{sy} - e^{ay}) \operatorname{Re}\left(\frac{\phi(s) - \widehat{\phi}(s)}{s^2}\right) ds \right| \\ &< \frac{1}{\pi} \int_{-\infty}^{\infty} e^{ay} (1 - \cos \omega y) \frac{C \|\mathbf{A}\|^4}{a^2 + \omega^2} d\omega \\ &= C \|\mathbf{A}\|^4 \frac{e^{ay} - 1}{a}. \end{aligned}$$

In the second line, we have used the fact that for all complex numbers x and y , $|\operatorname{Re}(x/y)| \leq |x|/|y|$. Thus, we have proved the $\|\mathbf{A}\|^4$ rate of convergence. The convergence rate of $\|\mathbf{A}\|^6$ can be shown in the same manner. ■

3.3. Numerical example 1. Our first example aims at illustrating the QTA procedure and demonstrating its accuracy in computing the loss distribution. In this example, we consider a hypothetical case where the correlation among obligors is moderate ($\|\mathbf{A}\|$ is small). For the more realistic case of strong correlation, and for the issue of market calibration, we defer to the next section, where we discuss extensions to the QTA method.

Consider a CDO structure with 125 names. We assume that the marginal default probability of each obligor is 2%, and the recovery rates range from 0% to 60%. Assume that obligors are arranged in ascending order according to their recovery rates (so that the first obligor has the highest loss given default and the last has the lowest). The correlation structure is

specified by the following 3-factor loading matrix:

$$(3.11) \quad \mathbf{A} = \begin{bmatrix} a & & \\ & b & \\ & & c \end{bmatrix}.$$

This factor structure is used to create inter- and intragroup correlations. For example, the parameter b represents the correlation between the high-loss group and the midloss group; it can be calibrated without affecting the low-loss obligors. (Similarly, calibrating c affects only correlation *within* the high-loss obligors.) This structure can be used to produce the base correlation skew observed in the market. For example, the base correlation curve for $a = .2$, $b = .5$, $c = .5$ is shown in Figure 1. (To plot this curve, we first use the 3-factor structure to price a hypothetical equity tranche whose attachment points are 0%– x %. The base correlation is then formally defined as the correlation parameter that the standard single-factor normal copula model requires to produce the same price. Figure 1 plots the base correlation against the upper attachment point x of the hypothetical equity tranche.)

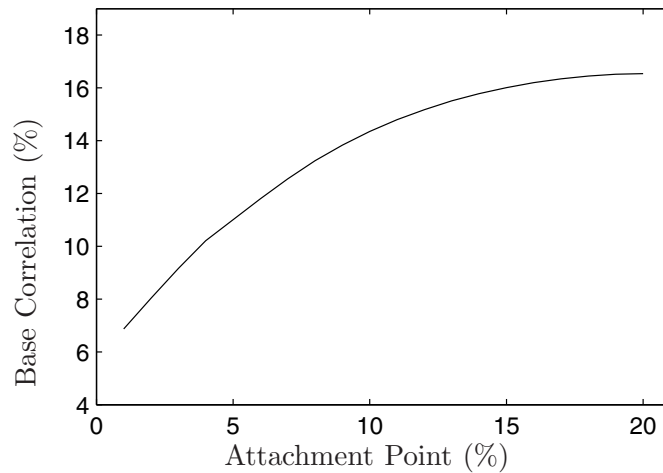


Figure 1. Skew produced by the 3-factor structure.

Let us now turn to the central issue of model computation. To obtain the loss distribution, we first compute the Laplace transform. The Laplace transform obtained by the QTA procedure, as compared to the “true” value (computed by Monte Carlo simulation), is shown in Figure 2.

The approximated Laplace transform can be used in the inversion formula to obtain the loss distribution and related expectations. For example, the expectation of the form (2.5) is shown in Figure 3.




Figure 2. Characteristic function.

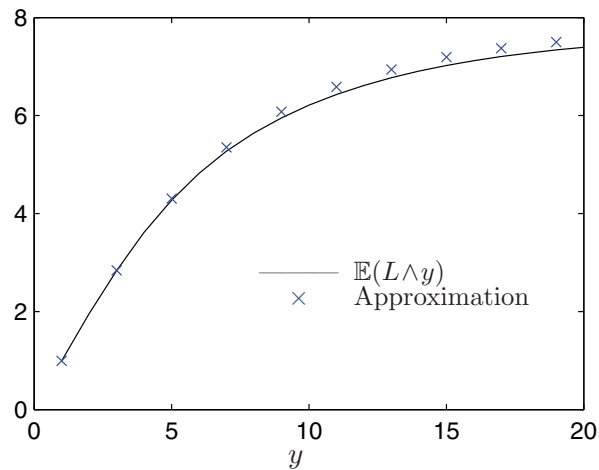


Figure 3. Inverting the approximated transform.

One can see that, when $\|\mathbf{A}\|$ is moderate, the approximation shows remarkable accuracy. For the case of stronger correlation, which can be observed in real market data, we must first discuss some methods for improving the accuracy. This will be the focus of the next section.

4. Improving the accuracy. When \mathbf{A} carries heavy loadings, the QTA method should be used in concurrence with either one of the methods described below. In the following sections, we assume without loss of generality that \mathbf{A} is structured so that its first column carries most of the loadings. The corresponding factor Z_1 thus affects most obligors, and so it will be referred to as the *market factor*.

4.1. Conditioning on the market factor. The d -factor normal copula model (2.3), once conditioned on the event $\{Z_1 = z\}$, becomes a normal copula model with $d-1$ factors whose loading matrix is lighter compared to that of the original d -factor model. Therefore, the QTA method of section 3 can be used to compute $\phi_z(s) := \mathbb{E}[e^{-sL} | Z_1 = z]$ for a given z . To obtain

the Laplace transform, we use one-dimensional numerical integration to evaluate the integral:

$$(4.1) \quad \phi(s) = \int_{-\infty}^{\infty} \phi_z(s) \varphi(z) dz,$$

where $\varphi(z) = e^{-z^2/2}/\sqrt{2\pi}$ is the standard normal density. In view of (4.1), the effort of computing the transform of a d -factor model is comparable to that of a single-factor model.

4.2. Segmenting the market factor. The QTA method relies on fitting $g_j(V_j)$ with a quadratic function (3.4). Goodness of fit can be improved by limiting the variation of the random variable V_j (defined in (3.3)). The method in this section reduces the variation of V_j by partitioning the market factor Z_1 into several *segments*. This approach, in contrast to that of the previous section, does not require the potentially time-consuming step of numerical integration.

Suppose that we segment Z_1 using the partition $-\infty = u_0 < u_1 < \dots < u_m = \infty$. Using (3.3), we write

$$(4.2) \quad \phi(s) = \mathbb{E} \left[e^{\sum_{j=1}^N \ln g_j(V_j)} \right] = \sum_{\ell=1}^m \mathbb{E} \left[e^{\sum_{j=1}^N \ln g_j(V_j)} \mathbf{1}_{\{u_{\ell-1} < Z_1 \leq u_{\ell}\}} \right].$$

By approximating the expectation for each segment separately, one obtains an estimate of $\phi(s)$. To approximate the expectation in a given segment, we posit an approximation of the form

$$(4.3) \quad \mathbb{E} \left[e^{\sum_{j=1}^N \ln g_j(V_j)} \mathbf{1}_{\{u < Z_1 \leq v\}} \right] \approx \mathbb{E} \left[e^{c + \mathbf{g}^T \mathbf{Z} + \mathbf{Z}^T \mathbf{H} \mathbf{Z}} \mathbf{1}_{\{u < Z_1 \leq v\}} \right],$$

where $u < v$ are real. As in section 3, this approximation is achieved by replacing $g_j(V_j)$ with a quadratic function $\alpha_j + \beta_j V_j + \eta_j V_j^2$ (see (3.4)), so that the exponent $\sum g_j(V_j)$ gets replaced by a quadratic function of \mathbf{Z} (see (3.5)). The resulting expectation (4.3) can ultimately be computed using a closed-form formula. (See Proposition 2.)

Segmenting Z_1 helps improve the approximation of $g_j(V_j)$ by making the distribution of V_j more concentrated. In particular, one can derive that

$$(4.4) \quad \mathbb{E}[V_j | u < Z_1 \leq v] = \frac{a_{j1}}{\sqrt{\mathbf{a}_j^T \mathbf{a}_j}} \frac{\varphi(v) - \varphi(u)}{\Phi(v) - \Phi(u)},$$

$$(4.5) \quad \text{Var}(V_j | u < Z_1 \leq v) = 1 - \frac{a_{j1}^2}{\mathbf{a}_j^T \mathbf{a}_j} \left[\left(\frac{\varphi(v) - \varphi(u)}{\Phi(v) - \Phi(u)} \right)^2 - \frac{\varphi'(v) - \varphi'(u)}{\Phi(v) - \Phi(u)} \right],$$

where $\varphi'(x) = -x \varphi(x)$. Of course, one can still use the least-squares method (3.7)–(3.8) to solve for coefficients in the approximation $g_j(V_j) \approx \alpha_j + \beta_j V_j + \eta_j V_j^2$, but the set of gridpoints Λ and the penalty function f used in the least-squares method should be chosen to reflect the conditional mean and variance (4.4)–(4.5). (That is, assuming that the segment $(u, v]$ is narrow so that V_j remains approximately normal, Λ should at least cover the range of ± 2 standard deviations around the mean, and $f(\lambda)$ should be exponentially decreasing in the square distance of λ from the mean.)

We now summarize the steps in the approximation (4.3):

1. For $j = 1, \dots, N$, compute $[\alpha_j \ \beta_j \ \eta_j]$ from (3.8), with Λ and f chosen to reflect the conditional mean and variance (4.4)–(4.5).
2. Compute c , \mathbf{g} , and \mathbf{H} from $[\alpha_j \ \beta_j \ \eta_j]$ using (3.6).
3. Compute the expectation on the right-hand side of (4.3) using the following proposition.

Proposition 2. *The expectation (4.3) simplifies to*

$$\frac{e^{c + \mathbf{g}^\top (\mathbf{I} - 2\mathbf{H})^{-1} \mathbf{g}/2}}{\sqrt{\det(\mathbf{I} - 2\mathbf{H})}} \left[\Phi\left(\frac{v - \nu}{\varsigma}\right) - \Phi\left(\frac{u - \nu}{\varsigma}\right) \right],$$

where ν is the first element of the vector $(\mathbf{I} - 2\mathbf{H})^{-1} \mathbf{g}$, and ς^2 is the top-left element of the square matrix $(\mathbf{I} - 2\mathbf{H})^{-1}$.

Note that the proposition involves computing the function Φ of complex numbers. We follow the definition in Abramowitz and Stegun [3, pp. 297, 932]:

$$(4.6) \quad \Phi(x) = \frac{1}{2} + \frac{x}{\sqrt{2\pi}} \sum_{k=0}^{\infty} \frac{(-x^2/2)^k}{k! (2k+1)}.$$

Computing $\Phi(x)$ for complex x is easy, since the series (4.6) converges rapidly.

4.3. Numerical example II. This example aims at demonstrating the accuracy of the two extended methods described in sections 4.1 and 4.2. We also investigate the issue of computing time for both methods.

With the methods of conditioning (section 4.1) or segmenting (section 4.2) the market factor, one can now handle the more realistic case of strong correlation between obligors. As a case study, let us consider the real market quotes at two different times: in September 2008 (when Lehman Brothers collapsed) and one year previously. The monthly fixings for the 5-year iTraxx Europe index in September 2007 and 2008 are shown in Table 1. (Source: Reuters.) The point is to contrast the shape of the base correlation curves implied by the quotes on these two days.

Table 1

Quoted tranche prices. Levels for the 0%–3% tranche are upfronts with a fixed 500bps spread and are quoted in terms of percentage of the notional principal. Levels for all other tranches are in basis points, with no fixed running spread.

	Tranches					Index
	0%–3%	3%–6%	6%–9%	9%–12%	12%–22%	
September 28, 2007	19%	92	37	24	15	36
September 28, 2008	42%	595	320	170	75	117

The base correlation curves implied from these quotes are shown in Figure 4.

Contrast, between the two days, the rate at which the base correlation increases from tranche to tranche. In September 2007 (the left plot), the base correlation increases at a progressively slower rate in senior tranches. In such cases, we find that a multifactor structure, like the one described in the example of section 3.3, can provide sufficient fit. On the other hand, in September 2008 (the right plot), the base correlation increases rapidly in the senior




Figure 4. Comparing the skew in September 2007 (left) and September 2008 (right).

tranches. This type of skew requires multifactor models that are extensions to the standard Gaussian copula model. We will fit the skew of September 2008 in section 5, where we develop approximation methods for multifactor models with stochastic loadings and factor-dependent recovery rates. For now, let us apply the methods of sections 4.1–4.2 to fit the data on September 28th, 2007.

We extend the 3-factor structure in the previous example (section 3.3) and use a 5-factor structure to fit the market data. (The idea is to calibrate the five factors in order to match the five tranches of the iTraxx index. We note that the number of factors does not adversely affect the computing time, since the QTA method bypasses the process of multidimensional numerical integration.) The 5-factor structure is shown below:

$$(4.7) \quad \mathbf{A} = \begin{bmatrix} \text{[Diagram of a 5-factor loading matrix A with 5 columns and varying row heights]} \end{bmatrix}.$$

The first factor is the market factor, the second and third factors control the correlation among the high-loss obligors, and the fourth and fifth factors control the correlation among the low-loss obligors. As before, we assume that each factor is homogeneous (in the sense that all nonzero elements in each column have the same value), so that there are five correlation parameters to calibrate.

The effect of each factor on the shape of the base correlation curve is consistent with our finding in the previous example: the market factor controls the level of the base correlations across all tranches, while the other factors control the skew of the curve. For example, we found, by manually calibrating the loadings, that the rough shape of the base correlation curve observed in September 2007 can be reproduced by letting the loading of the first factor

be 0.4 and the loading of the second factor be 0.9. We then use this as the initial starting point in an optimization procedure to minimize the sum square error between the market and model-implied base correlations.

The result of the fitted model is shown in Figure 5. We also show, in the same figure, the model-implied base correlation computed via Monte Carlo (MC) simulation. (Although this requires longer computing time compared to the QTA methods, the high number of replications allows us to obtain an almost-exact distribution of the loss in the 5-factor model. The corresponding 95% confidence interval lies within $\pm 1\%$ of the plotted value.)

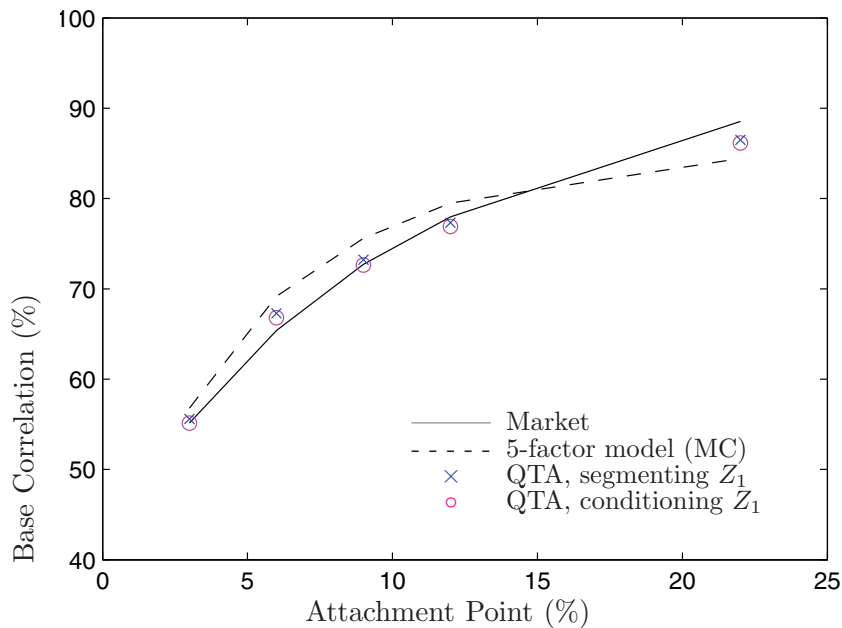


Figure 5. Result on September 28th, 2007.

As seen in Figure 5, the 5-factor structure succeeds in generating the general shape of the base correlation curve observed in the market, and the methods of sections 4.1 and 4.2 approximate the true 5-factor model quite well. Both methods also show remarkable speed compared to the traditional approach of computing the Laplace transform via multidimensional numerical integration, as shown in Table 2.

Table 2

Reduction in computing time, compared to the standard approach of numerical integration.

	Reduction in computing time
Conditioning	99.0%
Segmenting	99.3%

We end this example by emphasizing the issue of computing time. To compute the loss distribution via Laplace inversion formulae, one needs to compute $\phi(s)$ at *several* values of s . Furthermore, the calibration process requires recomputing the loss distribution at many different sets of parameters. This routine renders the standard approach of numerical integra-

tion infeasible for multifactor models. Similarly, Monte Carlo simulation is not appropriate for the process of calibration because of the large number of replications it requires to control the variance. Therefore, the reduction in computing time in Table 2 is crucial in the practical implementation of multifactor models.

5. Application to other models. The QTA method we have presented relies crucially on the assumption that the factors are normally distributed (so that Proposition 1 applies). So far we have worked under the standard normal copula model, in which the factor loadings \mathbf{A} are nonrandom and the losses-given-default Y_j are independent of the factors. The standard model, however, is incapable of modeling many important traits, such as dependence between recovery rate and defaults, and tail risks—a shortcoming that has led many researchers to consider factors with heavier tail distributions. Here, to stay within the class of normally distributed factor models, and thereby maintain the applicability of the QTA method, we consider some extensions of the normal copula model that address the weaknesses of the standard version (most of these extensions are proposed and analyzed in Andersen and Sidenius [5]). We also explain how the QTA method can be modified to deal with these extended models.

5.1. Correlation between defaults and recovery rates. Let the default times τ_j be specified, as usual, by (2.2)–(2.3). In the *random recovery (RR) model* considered in Andersen and Sidenius [5], dependence between defaults and recovery rates is introduced by letting the loss-given-default Y_j depend on the factors \mathbf{Z} :

$$(5.1) \quad Y_j = Y_j^{\max} \times [1 - R_j(\xi_j + U_j)], \quad \text{where } U_j := \mathbf{w}_j^\top \mathbf{Z}.$$

Here, $R_j : \mathbb{R} \rightarrow [0, 1]$ is a given increasing function, interpreted as the recoverable portion of the notional principal Y_j^{\max} . The ξ_j ($1 \leq j \leq N$) are independent random variables of prespecified distributions. The vector constant \mathbf{w}_j characterizes the dependence of Y_j on the economy-wide and/or group-specific factors \mathbf{Z} . Through \mathbf{w}_j , one can model intra- and intergroup correlations among recovery rates, just as one models the correlations among defaults through the loading matrix \mathbf{A} . Note that, in this setting, τ_j and Y_j remain independent conditional on \mathbf{Z} . This allows us to apply the same idea in earlier sections to approximate the Laplace transform.

As before, we first compute the conditional transform $\mathbb{E}[e^{-sL} | \mathbf{Z}]$. Then we let $\psi_j(s, u) := \mathbb{E}[e^{-sY_j} | U_j = u]$ be the Laplace transform of Y_j conditioned on U_j . Mimicking the steps in (2.9) of section 2.2, we now yield

$$\mathbb{E}[e^{-sL} | \mathbf{Z}] = \prod_{j=1}^N \left(1 + (\psi_j(s, U_j) - 1) \Phi \left(\frac{x_j - \mathbf{a}_j^\top \mathbf{Z}}{\sqrt{1 - \mathbf{a}_j^\top \mathbf{a}_j}} \right) \right).$$

Thus, similar to (3.3), we now have

$$(5.2) \quad \mathbb{E}[e^{-sL} | \mathbf{Z}] = \prod_{j=1}^N e^{g_j(V_j, U_j)} = e^{\sum_{j=1}^N g_j(V_j, U_j)},$$

where g_j follows the old definition (3.2), but with $\psi_j(s)$ replaced by $\psi_j(s, U_j)$. The next step is to approximate g_j by a quadratic function. Instead of approximating g_j as a function of V_j

alone (as we did in (3.4) of section 3.1), we now fit

$$(5.3) \quad g_j(V_j, U_j) \approx \alpha_j + \beta_j V_j + \kappa_j U_j + \eta_j V_j^2 + \zeta_j U_j V_j + \nu_j U_j^2,$$

where the coefficients $(\alpha_j, \beta_j, \kappa_j, \eta_j, \zeta_j, \nu_j)$ are obtained by the least-squares method described in section 3.1. Because V_j and U_j are linear in \mathbf{Z} , the approximant (5.3) simplifies to a quadratic function of \mathbf{Z} . Consequently,

$$(5.4) \quad \phi(s) = \mathbb{E} \left[\mathbb{E} [e^{-sL} | \mathbf{Z}] \right] = \mathbb{E} \left[e^{\sum_{j=1}^N g_j(V_j, U_j)} \right] \approx \mathbb{E} \left[e^{c + \mathbf{g}^\top \mathbf{Z} + \mathbf{Z}^\top \mathbf{H} \mathbf{Z}} \right],$$

where the scalar c , the vector \mathbf{g} , and the matrix \mathbf{H} are derived from the coefficients of (5.3). Proposition 1 can then be evoked to evaluate the expectation.

5.2. QTA on multifactor models with stochastic factor loadings. The QTA method can accommodate certain types of multifactor models with stochastic factor loadings. Two such models are discussed here, each with the instruction on how to apply the QTA method.

Schloegl [20], as well as Burtschell, Gregory, and Laurent [7, 6], randomize the factor loadings by “mixing copulae.” The simplest version of such an approach is described in its multifactor form as follows. Let the default times be given, as before, by $\tau_j = F_j^{-1}(\Phi(X_j))$ (see (2.2)). But instead of specifying X_j as in (2.3), we now let

$$(5.5) \quad X_j = B_j \left(\mathbf{a}_j^\top \mathbf{Z} + \sqrt{1 - \mathbf{a}_j^\top \mathbf{a}_j} \epsilon_i \right) + (1 - B_j) \left(\tilde{\mathbf{a}}_j^\top \mathbf{Z} + \sqrt{1 - \tilde{\mathbf{a}}_j^\top \tilde{\mathbf{a}}_j} \epsilon_i \right),$$

where $\mathbf{a}_j, \tilde{\mathbf{a}}_j$ are constant vectors, and B_j are independent Bernoulli random variables. While the X_j in (5.5) remains standard normal, it is now a “mixture” of two copulae: if $B_j = 1$, then the copula \mathbf{a}_j prevails; if $B_j = 0$, then $\tilde{\mathbf{a}}_j$ prevails.

To apply the QTA method, we begin, as usual, by computing the conditional transform $\mathbb{E} [e^{-sL} | \mathbf{Z}]$. Define

$$(5.6) \quad V_j := -\frac{\mathbf{a}_j^\top \mathbf{Z}}{\sqrt{\mathbf{a}_j^\top \mathbf{a}_j}} \quad \text{and} \quad \tilde{V}_j := -\frac{\tilde{\mathbf{a}}_j^\top \mathbf{Z}}{\sqrt{\tilde{\mathbf{a}}_j^\top \tilde{\mathbf{a}}_j}}.$$

By following the same line of arguments that leads to (3.3), one now obtains

$$\mathbb{E} [e^{-sL} | \mathbf{Z}] = \prod_{j=1}^N \left[r_j g_j(V_j) + (1 - r_j) \tilde{g}_j(\tilde{V}_j) \right] = e^{\sum_{j=1}^N \ln [r_j g_j(V_j) + (1 - r_j) \tilde{g}_j(\tilde{V}_j)]},$$

where \tilde{g}_j is defined similarly to g_j (see (3.2)), only with $\tilde{\mathbf{a}}_j$ replacing \mathbf{a}_j , and $r_j = P(B_j = 1)$. Note the similarity between the above equation and (3.3); the g_j in (3.3) is now replaced by a *mixture* of g_j and \tilde{g}_j . The remaining steps are similar to (5.3)–(5.4). Using the least-squares method, we approximate

$$(5.7) \quad \ln \left[r_j g_j(V_j) + (1 - r_j) \tilde{g}_j(\tilde{V}_j) \right] \approx \text{degree-2 polynomial in } V_j \text{ and } \tilde{V}_j.$$

This polynomial simplifies to a quadratic function of \mathbf{Z} . Consequently,

$$\phi(s) = \mathbb{E} \left[e^{\sum_{j=1}^N \ln [r_j g_j(V_j) + (1 - r_j) \tilde{g}_j(\tilde{V}_j)]} \right] \approx \mathbb{E} \left[e^{c + \mathbf{g}^\top \mathbf{Z} + \mathbf{Z}^\top \mathbf{H} \mathbf{Z}} \right].$$

Proposition 1 can then be evoked to evaluate this expectation. Note that the above method is also applicable to the “three-stage” model considered in Burtschell, Gregory, and Laurent [7], which can be viewed as a variation of the mixed copula model.

As another method of randomizing the factor loadings, the *random factor loading (RFL) model*, proposed and analyzed in Andersen and Sidenius [5], specifies \mathbf{A} as a function of \mathbf{Z} . Let us consider a multifactor case where the factor loadings depend only on the market factor Z_1 :

$$(5.8) \quad X_j = \mathbf{a}_j^\top \mathbf{Z} + b_j \epsilon_j + \xi_j, \quad \mathbf{a}_j = \begin{cases} \mathbf{a}_j^{(\text{high})} & \text{if } Z_1 \leq \mu, \\ \mathbf{a}_j^{(\text{low})} & \text{if } Z_1 > \mu, \end{cases}$$

where $\mathbf{a}_j^{(\text{high})}$ and $\mathbf{a}_j^{(\text{low})}$ are constant vectors. The constants b_j and ξ_j are chosen such that X_j has mean 0 and variance 1. The intuition of the model is that the loadings depend on the market condition (Z_1), allowing for stronger correlation in a bear market. Note that, here, X_j is no longer Gaussian. Therefore, (2.2) should be changed to $\tau_j = F_j^{-1}(\Theta_j(X_j))$, where Θ_j is the cumulative distribution function (cdf) of X_j . (The expression (2.8) for the conditional default probability should likewise be adjusted to reflect the distribution of X_j . In particular, the definition of x_j in (2.8) should be changed to $x_j = \Theta_j^{-1}(F_j(t)) - \xi_j$.)

To compute the Laplace transform in this setting, we apply the idea of segmenting the market factor from section 4.2. More specifically, the approximation takes the form

$$(5.9) \quad \phi(s) \approx \sum_{\ell=1}^m \mathbb{E} \left[e^{c_\ell + \mathbf{g}_\ell^\top \mathbf{Z} + \mathbf{Z}^\top \mathbf{H}_\ell \mathbf{Z}} 1_{\{u_{\ell-1} < Z_1 \leq u_\ell\}} \right].$$

For each segment, b_ℓ , \mathbf{g}_ℓ , and \mathbf{H}_ℓ are computed using the analytical method described in section 4.2; we use the loading $\mathbf{a}_j^{(\text{high})}$ for the segments that belong to the region $\{Z_1 \leq \mu\}$ and $\mathbf{a}_j^{(\text{low})}$ for those belonging to $\{Z_1 > \mu\}$. Since each expectation in (5.9) is taken over the Gaussian factors \mathbf{Z} , Proposition 2 still applies to evaluate them in closed form.

5.3. Numerical example III. We now test the QTA method on the extended Gaussian copula models with either stochastic recovery rates or stochastic factor loadings. With these extensions of the standard Gaussian copula model, we are now able to consider the data on September 28th, 2008 (shown in the example of section 4.3), in which correlation among reference entities is high, and the base correlation curve exhibits a very pronounced skew (see Figure 4). Andersen and Sidenius [5] have already examined the suitability of the RR model and the RFL model in fitting such a skew. The main objective of this example is to examine the benefit of including more factors in these models. More importantly, we will demonstrate the accuracy of our approximation method as it applies to the multifactor version of these models.

We first consider the RR model in section 5.1. Again, we use the 5-factor structure with group-specific factors similar to that of section 4.3, but now we introduce dependence between defaults and recovery rates using (5.1), with R_j chosen to be the standard Gaussian cdf Φ (see, for example, Andersen and Sidenius [5]). The notional principle of each obligor Y_j^{max} is assumed to be one for all j . For simplicity, we assume that the recovery rates depend on the factors only through the market factor Z_1 , i.e., $U_j = w Z_1$ for all j (so that, conditioned on Z_1 , the model becomes a 4-factor standard Gaussian copula model).

The guideline for calibrating the factor loadings is similar to the example in section 4.3; here, we find that the added correlation between defaults and recovery rates (i.e., the parameter w) helps lift the base correlation in the senior tranche. By trial and error, we find that the rough shape of the base correlation curve can be generated by letting the loading of the first factor be 0.4, letting the loading of the second factor be 0.7, and letting $w = 0.9$. We use this set of parameters as the starting point in an optimization procedure to minimize the sum square error between the market and model-implied base correlations. The base correlation of the fitted model is shown in Figure 6.




Figure 6. Model with correlated defaults and recovery rates.

Evidently, the 5-factor RR model fits the market almost perfectly. To understand the inherent multifactor nature of the market on this day, we perform principal component analysis (PCA) on the correlation structure among defaults and recovery rates, so as to extract the closest single-factor structure from the 5-factor structure. The single-factor structure extracted from the PCA procedure is shown in the same plot. As seen in Figure 6, the single-factor RR model is capable of producing a very steep base correlation curve, but allowing a multifactor structure serves as a way of fine tuning the model to match, almost exactly, the market quotes. With the QTA method, the 5-factor model can be handled at almost the same computational cost as the single-factor model, since the QTA method does not involve multidimensional numerical integration.

As the last case study, let us use a multifactor RFL model (5.8) to fit the same data and use the QTA method to compute the model output.

The correlation structure in the multifactor RFL model is as follows. We assume that the 5-factor structure (4.7) prevails in the region $Z_1 > \mu$. For $Z_1 \leq \mu$, we assume, for simplicity, that a single-factor structure (which loads only on Z_1) prevails. As before, we assume that each factor is homogeneous, in the sense that it carries the same loading for every obligor that it affects. For comparison, we will also consider a homogeneous, single-factor RFL model.

In the single-factor RFL model, there are three correlation parameters to calibrate: μ , $a^{(\text{high})}$, and $a^{(\text{low})}$. (For a comprehensive discussion of the effect these parameters have on the correlation skew, we refer the reader to Andersen and Sidenius [5].) Figure 7 shows the base correlation skew when $\mu = -1.7$, $a^{(\text{low})} = .6$, and $a^{(\text{high})} = .99$. We find that this set of parameters produces the base correlation curve that has the steepest upward trend. Nevertheless, it cannot match the steepness of the market-implied base correlation curve. Readjusting the parameters in the single-factor model will not solve the problem; doing so will, at best, result in a parallel shift of the base correlation curve.

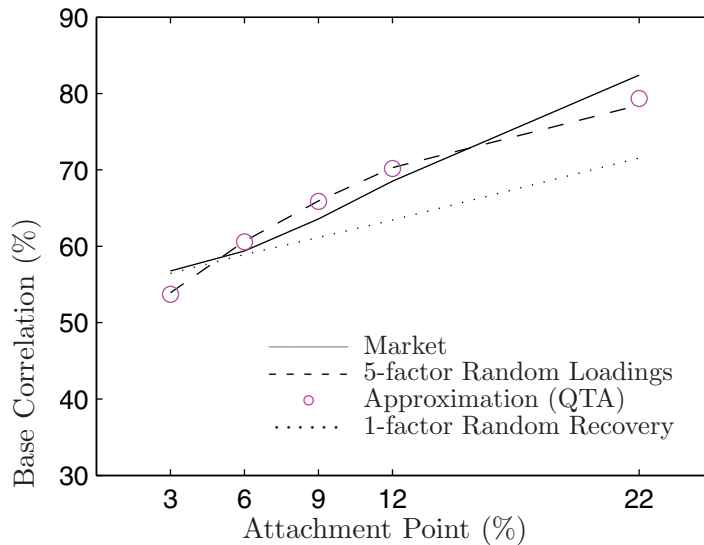


Figure 7. Stochastic factor loading.

The multifactor structure (4.7) provides a way to control nonparallel shifting of the base correlation curve. The second factor, for example, is specific to the high-loss obligors and can be calibrated without affecting the rest of the obligors in the portfolio. In our calibration, we find that the second factor, which carries the loading of 0.8 in our fitted model, plays an important role in steepening the base correlation curve, so that it matches the overall shape of the market-implied curve (see Figure 7). As for the central issue of the QTA method’s accuracy, one can see that the method approximates the 5-factor RFL model quite well.

6. Conclusion. The multifactor version of the Gaussian copula models, as compared to its single-factor counterpart, offers more richness in modeling the correlation structure of defaults and is useful in generating the base correlation skew observed in the market. Indeed, even models that are extensions of the standard Gaussian copula model, such as those with stochastic recovery rates and stochastic factor loadings, are shown in our examples to benefit from the inclusion of more factors. However, computing the loss distribution in multifactor models becomes more challenging as the number of factors increases. This paper proposes an analytical method for approximating the loss distribution through a closed-form approximation of its Laplace transform. The relevance of the approach is that it solves the curse of dimensionality that is usually associated with multifactor models. The method is fast,

accurate, and applicable to a range of models with Gaussian factors. Our numerical examples provide a case study of the viability of the method.

Appendices: Proofs.

Appendix A. Proof of Propositions 1 and 2. The following is the proof of Proposition 2. The proof of Proposition 1 follows the same line of argument.

First, assume that c , \mathbf{g} , and \mathbf{H} have no imaginary parts. Let $\mathbf{V} := \mathbf{I} - 2\mathbf{H}$, and let $\mathbf{h} := \mathbf{V}^{-1}\mathbf{g}$:

$$\begin{aligned} & \mathbb{E} \left[e^{c + \mathbf{g}^\top \mathbf{Z} + \mathbf{Z}^\top \mathbf{H} \mathbf{Z}} 1\{u < Z_1 \leq v\} \right] \\ &= \int \dots \int \frac{1}{\sqrt{(2\pi)^d}} 1\{u < Z_1 \leq v\} e^{c + \mathbf{g}^\top \mathbf{Z} + \mathbf{Z}^\top \mathbf{H} \mathbf{Z}} e^{-\mathbf{Z}^\top \mathbf{Z}/2} dZ_1 \dots dZ_d \\ &= \int \dots \int \frac{e^{c + \mathbf{g}^\top \mathbf{h}/2}}{\sqrt{(2\pi)^d}} 1\{u < Z_1 \leq v\} e^{-\frac{1}{2}(\mathbf{Z} - \mathbf{h})^\top \mathbf{V} (\mathbf{Z} - \mathbf{h})} dZ_1 \dots dZ_d \\ &= \frac{e^{c + \mathbf{g}^\top \mathbf{h}/2}}{\sqrt{\det \mathbf{V}}} \left[\Phi\left(\frac{v - \nu}{\varsigma}\right) - \Phi\left(\frac{u - \nu}{\varsigma}\right) \right], \end{aligned}$$

where ν is the first element of the vector \mathbf{h} , and ς^2 is the top-left element of the matrix \mathbf{V}^{-1} . The last equation follows from the fact that the last integral is taken with respect to the multidimensional Gaussian distribution with mean \mathbf{h} and covariance matrix \mathbf{V}^{-1} . The above derivation is still valid when c , \mathbf{g} , and \mathbf{H} are complex, provided that the Gaussian cdf Φ is interpreted as the power series (4.6).

Appendix B. Proof of Theorem 1. The following notations are useful:

$$h_j(\varepsilon, \lambda, \mathbf{A}) := \frac{x_j + \lambda \sqrt{\mathbf{a}_j^\top \mathbf{a}_j}}{\sqrt{1 - \varepsilon^2 \mathbf{a}_j^\top \mathbf{a}_j}},$$

where \mathbf{a}_j^\top is the j th row of \mathbf{A} . Define

$$(B.1) \quad G_j(s, x) := 1 + (\psi_j(s) - 1)\Phi(x), \quad G_j^{(k)}(s, x) := \frac{\partial^k}{\partial x^k} G_j(s, x).$$

Suppose that $\Lambda = \{\lambda_1, \dots, \lambda_r\}$. Define

$$\mathbf{L}^{(k)} := \begin{bmatrix} 1 & \lambda_1 & \lambda_1^2 & \dots & \lambda_1^k \\ \vdots & & & & \vdots \\ 1 & \lambda_r & \lambda_r^2 & \dots & \lambda_r^k \end{bmatrix}, \quad \mathbf{\Gamma}_j^{(k)}(\varepsilon, s, \mathbf{A}) := \frac{\partial^k}{\partial \varepsilon^k} \begin{bmatrix} \ln G_j(s, h(\varepsilon, \varepsilon \lambda_1, \mathbf{A})) \\ \vdots \\ \ln G_j(s, h(\varepsilon, \varepsilon \lambda_r, \mathbf{A})) \end{bmatrix}.$$

Lemma 1. *There exist a positive number $\varepsilon^* < 1$ and a finite number C such that for all $j = 1, \dots, N$ and for all $k = 1, \dots, 8$,*

$$(B.2) \quad \left| \frac{\partial^k}{\partial \varepsilon^k} \ln G_j(s, h_j(\varepsilon, \varepsilon \lambda, \mathbf{A})) \right| < C \text{ whenever } |\varepsilon| < \varepsilon^*, s \in \mathbb{C}^+, \lambda \in \Lambda, \|\mathbf{A}\| \leq 1.$$

Furthermore, for fixed s and \mathbf{A} , there exist $\gamma_{jkl} \in \mathbb{C}$ such that for all $\lambda \in \mathbb{R}$,

$$(B.3) \quad \frac{\partial^k}{\partial \varepsilon^k} \Big|_{\varepsilon=0} \ln G_j(s, h_j(\varepsilon, \varepsilon\lambda, \mathbf{A})) = \sum_{\ell=0}^k \lambda^\ell \gamma_{jkl},$$

where $\gamma_{jkl} = 0$ when $k - \ell$ is odd. In particular, $\mathbf{\Gamma}_j^{(k)}(0, s) = \mathbf{L}^{(k)} [\gamma_{jk0} \ \gamma_{jk1} \ \dots \ \gamma_{jkk}]^\top$.

Proof. It can be shown that $|\Phi^{(k)}(x)| < 15$ for all $x \in \mathbb{R}$ and $k \leq 8$. Note also that $|\psi_j(s)| \leq 1$. Pick a negative number $x^* > \max_j x_j$. From the definition of G_j , one has

$$(B.4) \quad 1 - 2\Phi(x^*) < |G_j(s, x)| \leq 1 \quad \text{and} \quad |G_j^{(k)}(s, x)| < 31 \quad \forall x < x^*, s \in \mathbb{C}^+, k = 0, 1, \dots, 8.$$

Pick $\varepsilon^* < 1$ such that $h_j(\varepsilon, \varepsilon\lambda, \mathbf{A}) < x^*$ whenever $|\varepsilon| < \varepsilon^*$, $\lambda \in \Lambda$, $\|\mathbf{A}\| \leq 1$. It can be shown that

$$(B.5) \quad \left| \frac{\partial^k}{\partial \varepsilon^k} h_j(\varepsilon, \varepsilon\lambda, \mathbf{A}) \right| < \frac{(2k)! (|x_j| + 2|\lambda|)}{(1 - \varepsilon^* 2)^{k+1/2}} \quad \forall \varepsilon < \varepsilon^*, \|\mathbf{A}\| \leq 1.$$

Property (B.2) follows from (B.4) and (B.5).

Equation (B.3) is a result of the following observation. Given any mapping $(\varepsilon, x) \mapsto h(\varepsilon, x)$, the derivative $(\partial^k / \partial \varepsilon^k)|_{\varepsilon=0} h(\varepsilon, \varepsilon\lambda)$, if it exists, can be shown to take the form of a degree- k polynomial in λ , with the coefficient of λ^ℓ ($\ell \leq k$) identified as

$$(B.6) \quad \frac{1}{\ell!} \frac{\partial^{k+\ell}}{\partial \varepsilon^k \partial \lambda^\ell} \Big|_{\varepsilon=0} h(\varepsilon, \varepsilon\lambda) = \frac{1}{\ell!} \frac{\partial^{k+\ell}}{\partial \varepsilon^k \partial \lambda^\ell} \Big|_{\varepsilon=0} \varepsilon^\ell h(\varepsilon, \lambda) = \binom{k}{\ell} \frac{\partial^k}{\partial \lambda^\ell \partial \varepsilon^{k-\ell}} \Big|_{\varepsilon=0} h(\varepsilon, \lambda).$$

Furthermore, if $h(\varepsilon, \lambda) = h(-\varepsilon, \lambda)$ (even function), then the derivative in (B.6) is zero whenever $k - \ell$ is odd. (This implies that γ_{jkl} in (B.3) is zero whenever $k - \ell$ is odd.) ■

Let $\mathbf{W} = \text{diag}(f(\lambda_1), \dots, f(\lambda_r))$. Let $\mathbf{b}_0, \mathbf{b}_1$, and \mathbf{b}_2 be the first, second, and third rows of the $3 \times r$ matrix $(\mathbf{L}^{(2)\top} \mathbf{W} \mathbf{L}^{(2)})^{-1} \mathbf{L}^{(2)\top} \mathbf{W}$. Define

$$(B.7) \quad S_{\mathbf{A}}^{(k)}(\varepsilon, s, \mathbf{Z}) := \sum_{j=1}^N \frac{\partial^k}{\partial \varepsilon^k} \ln G_j(s, h(\varepsilon, \varepsilon V_j, \mathbf{A})),$$

$$(B.8) \quad \widehat{S}_{\mathbf{A}}^{(k)}(\varepsilon, s, \mathbf{Z}) := \sum_{j=1}^N (\mathbf{b}_0 + V_j \mathbf{b}_1 + V_j^2 \mathbf{b}_2) \mathbf{\Gamma}_j^{(k)}(\varepsilon, s, \mathbf{A}) = \frac{\partial^k}{\partial \varepsilon^k} \widehat{S}_{\mathbf{A}}^{(0)}(\varepsilon, s, \mathbf{Z}).$$

It follows that

$$(B.9) \quad \phi_{\varepsilon \mathbf{A}}(s) = \mathbb{E} \left[e^{S_{\mathbf{A}}^{(0)}(\varepsilon, s, \mathbf{Z})} \right], \quad \widehat{\phi}_{\varepsilon \mathbf{A}}(s) = \mathbb{E} \left[e^{\widehat{S}_{\mathbf{A}}^{(0)}(\varepsilon, s, \mathbf{Z})} \right].$$

This equation is best understood when $\varepsilon = 1$. When $\varepsilon = 1$, $S_{\mathbf{A}}^{(0)}(\varepsilon, s, \mathbf{Z})$ is the same as $\sum_j \ln g_j(V_j)$ in (3.3); thus the first part of (B.9) follows. To see the second part, note that (3.8) is equivalent to $\alpha_j = \mathbf{b}_0 \mathbf{\Gamma}_j(\varepsilon, s, \mathbf{A})$, $\beta_j = \mathbf{b}_1 \mathbf{\Gamma}_j(\varepsilon, s, \mathbf{A})$, and $\eta_j = \mathbf{b}_2 \mathbf{\Gamma}_j(\varepsilon, s, \mathbf{A})$. Therefore, $\widehat{S}_{\mathbf{A}}^{(0)}(\varepsilon, s, \mathbf{Z}) = \sum_j (\alpha_j + \beta_j V_j + \eta_j V_j^2)$. Thus the second part of (B.9) follows.

Proposition 3. *There exist $\varepsilon_0 > 0$ and integrable functions $g(\mathbf{Z}), \widehat{g}(\mathbf{Z})$ such that*

$$(B.10) \quad \left| \frac{\partial^k}{\partial \varepsilon^k} e^{S_{\mathbf{A}}^{(0)}(\varepsilon, s, \mathbf{Z})} \right| < g(\mathbf{Z}) \quad \text{and} \quad \left| \frac{\partial^k}{\partial \varepsilon^k} e^{\widehat{S}_{\mathbf{A}}^{(0)}(\varepsilon, s, \mathbf{Z})} \right| < \widehat{g}(\mathbf{Z})$$

for all $|\varepsilon| < \varepsilon_0$, $s \in \mathbb{C}^+$, $k \leq 8$ and for all \mathbf{A} such that $\|\mathbf{A}\| \leq 1$.

Proof. From (B.7),

$$\frac{\partial^k}{\partial \varepsilon^k} e^{S_{\mathbf{A}}^{(0)}(\varepsilon, s, \mathbf{Z})} = \frac{\partial^k}{\partial \varepsilon^k} \prod_{j=1}^N G_j(s, h(\varepsilon, \varepsilon V_j, \mathbf{A})).$$

Using (B.4) and (B.5), the above expression can be bounded by a polynomial in $|V_j|$; thus the first part of (B.10) follows.

For any $G(\varepsilon)$ twice continuously differentiable, there exists $\varepsilon' \in (0, \varepsilon)$ such that $G(\varepsilon) = G(0) + \varepsilon G^{(1)}(0) + \frac{1}{2} \varepsilon^2 G^{(2)}(\varepsilon')$ (mean value theorem). Similarly, Lemma 1 implies that for every $\varepsilon < \varepsilon^*$, there exists $\varepsilon' \in (0, \varepsilon)$ such that

$$(B.11) \quad \ln G_j(s, h(\varepsilon, \varepsilon \lambda, \mathbf{A})) = \gamma_{j00} + \varepsilon \gamma_{j11} \lambda + \frac{1}{2} \varepsilon^2 \frac{\partial^2}{\partial \varepsilon^2} \Big|_{\varepsilon=\varepsilon'} \ln G_j(s, h(\varepsilon, \varepsilon \lambda, \mathbf{A})).$$

Note that the second derivative in the last term is bounded by the constant C in (B.2). In vector form, $\mathbf{\Gamma}_j^{(0)}(\varepsilon, s, \mathbf{A}) = \mathbf{L}^{(1)}[\gamma_{j00} \quad \varepsilon \gamma_{j11}]^\top + \frac{1}{2} \varepsilon^2 \mathbf{\Xi}_j$ for some column vector $\mathbf{\Xi}_j$ with $\|\mathbf{\Xi}_j\| < C$. Substituting $\mathbf{\Gamma}_j^{(0)}(\varepsilon, s, \mathbf{A})$ into (B.8), noting that $\mathbf{b}_2 \mathbf{L}^{(1)} = \mathbf{0}$, one yields

$$(B.12) \quad \left| \widehat{S}_{\mathbf{A}}^{(0)}(\varepsilon, s, \mathbf{Z}) \right| \leq \sum_{j=1}^N \left| (\mathbf{b}_0 + V_j \mathbf{b}_1) \mathbf{\Gamma}_j(\varepsilon, s, \mathbf{A}) + \frac{\varepsilon^2 \mathbf{b}_2 \mathbf{\Xi}_j V_j^2}{2} \right| \quad \forall \varepsilon < \varepsilon^*.$$

Let $\|\mathbf{b}_0\|$, $\|\mathbf{b}_1\|$, and $\|\mathbf{b}_2\|$ be bounded by B . Noting that $|V_j| \leq \sqrt{\mathbf{Z}^\top \mathbf{Z}}$ (see (3.3)), and $|\mathbf{\Gamma}_j(\cdot)| < C$, it follows from (B.12) that

$$(B.13) \quad \left| e^{\widehat{S}_{\mathbf{A}}^{(0)}(\varepsilon, s, \mathbf{Z})} \right| \leq e^{\left| \widehat{S}_{\mathbf{A}}^{(0)}(\varepsilon, s, \mathbf{Z}) \right|} < e^{NBC(1 + \sqrt{\mathbf{Z}^\top \mathbf{Z}} + \varepsilon^2 \mathbf{Z}^\top \mathbf{Z}/2)}$$

for all $\varepsilon < \varepsilon^*$ and $s \in \mathbb{C}^+$. Similarly, we can bound (B.8) by

$$(B.14) \quad \left| \widehat{S}_{\mathbf{A}}^{(k)}(\varepsilon, s, \mathbf{Z}) \right| \leq NBC \left(1 + \sqrt{\mathbf{Z}^\top \mathbf{Z}} + \mathbf{Z}^\top \mathbf{Z} \right).$$

Choose $\varepsilon_0 < \min\{\varepsilon^*, 1/\sqrt{NBC}\}$. The second part of (B.10) then follows from (B.14) and (B.13) (chain rule). \blacksquare

Lemma 2. *Suppose that $\|\mathbf{A}\| \leq 1$. If $\overline{\lambda^1} = \overline{\lambda^3} = \overline{\lambda^5} = 0$, then for $n = 0, 1, 2, 3$,*

$$(B.15) \quad \frac{\partial^n}{\partial \varepsilon^n} \Big|_{\varepsilon=0} \frac{\phi_{\varepsilon \mathbf{A}}(s) - \widehat{\phi}_{\varepsilon \mathbf{A}}(s)}{\phi_0(s)} = \mathbb{E} \left[\frac{\partial^n}{\partial \varepsilon^n} \Big|_{\varepsilon=0} \frac{e^{S_{\mathbf{A}}^{(0)}(\varepsilon, s, \mathbf{Z})} - e^{\widehat{S}_{\mathbf{A}}^{(0)}(\varepsilon, s, \mathbf{Z})}}{\phi_0(s)} \right] = 0.$$

If, in addition, $\overline{\lambda^2} = 1$, $\overline{\lambda^4} = 3$, and $\overline{\lambda^6} = 15$, then (B.15) holds for $n = 4, 5$ as well.

Proof. If $\overline{\lambda^1} = \overline{\lambda^3} = \overline{\lambda^5} = 0$, then $\mathbf{b}_0\mathbf{L}^{(3)} = [1 \ 0 \ 0 \ 0]$, $\mathbf{b}_1\mathbf{L}^{(3)} = [0 \ 1 \ 0 \ \overline{\lambda^4}/\overline{\lambda^2}]$, and $\mathbf{b}_2\mathbf{L}^{(3)} = [0 \ 0 \ 1 \ 0]$. Let $r_k := S_{\mathbf{A}}^{(k)}(0, s, \mathbf{Z})$ and $\widehat{r}_k := \widehat{S}_{\mathbf{A}}^{(k)}(0, s, \mathbf{Z})$. Using Lemma 1, one yields

$$r_1 = \widehat{r}_1 = \sum_{j=1}^N \gamma_{j11}V_j, \quad r_2 = \widehat{r}_2 = \sum_{j=1}^N (\gamma_{j22}V_j^2 + \gamma_{j20}), \quad r_3 - \widehat{r}_3 = \sum_{j=1}^N \gamma_{j33} \left(V_j^3 - V_j \frac{\overline{\lambda^4}}{\overline{\lambda^2}} \right).$$

Note that $\mathbb{E}[r_3 - \widehat{r}_3] = 0$. The derivative inside the expectation (B.15), once written in terms of r_k and \widehat{r}_k ($k \leq n$) using the chain rule, simplifies to zero for $n = 0, 1, 2$ and to $r_3 - \widehat{r}_3$ for $n = 3$. Obviously, (B.15) holds for $n = 0, 1, 2, 3$.

If, in addition, $\overline{\lambda^2} = 1$, $\overline{\lambda^4} = 3$, and $\overline{\lambda^6} = 15$, then $\mathbf{b}_0\mathbf{L}^{(5)} = [1 \ 0 \ 0 \ 0 \ -3 \ 0]$, $\mathbf{b}_1\mathbf{L}^{(5)} = [0 \ 1 \ 0 \ 3 \ 0 \ 15]$, and $\mathbf{b}_2\mathbf{L}^{(5)} = [0 \ 0 \ 1 \ 0 \ 6 \ 0]$. Using Lemma 1, one obtains

$$r_4 - \widehat{r}_4 = \sum_{j=1}^N \gamma_{j44}(V_j^4 - 6V_j^2 + 3), \quad r_5 - \widehat{r}_5 = \sum_{j=1}^N [\gamma_{j55}(V_j^5 - 15V_j) + \gamma_{j53}(V_j^3 - 3V_j)].$$

For $n = 4$, the inside of (B.15) simplifies to $r_4 - \widehat{r}_4 + 4r_1(r_3 - \widehat{r}_3)$. Using the fact that $\mathbb{E}[V_j V_k^3] = 3\mathbb{E}[V_j V_k]$ for all j, k (property of $\mathcal{N}(0, 1)$ random variables), one can show that $\mathbb{E}[r_4 - \widehat{r}_4] = 0$ and $\mathbb{E}[r_1(r_3 - \widehat{r}_3)] = 0$. Consequently, (B.15) holds for $n = 4$. For $n = 5$, the expectation (B.15) simplifies to $\mathbb{E}[r_5 - \widehat{r}_5] + 5\mathbb{E}[r_1(r_4 - \widehat{r}_4)] + 10\mathbb{E}[(r_3 - \widehat{r}_3)(r_2 - r_1^2)]$. These three expectations can be easily shown to be zero. Thus, (B.15) holds for $n = 5$. ■

We are now ready to prove Theorem 1 in section 3.2. From the mean value theorem,

$$\phi_{\varepsilon\mathbf{A}}(s) - \widehat{\phi}_{\varepsilon\mathbf{A}}(s) = \sum_{k=0}^{n-1} \frac{\varepsilon^k}{k!} \frac{\partial^k}{\partial \varepsilon^k} \Big|_{\varepsilon=0} [\phi_{\varepsilon\mathbf{A}}(s) - \widehat{\phi}_{\varepsilon\mathbf{A}}(s)] + \frac{\varepsilon^n}{n!} \frac{\partial^n}{\partial \varepsilon^n} \Big|_{\varepsilon=\varepsilon'} [\phi_{\varepsilon\mathbf{A}}(s) - \widehat{\phi}_{\varepsilon\mathbf{A}}(s)].$$

Proposition 3 implies that the derivative in the last term is bounded by $\mathbb{E}[g(\mathbf{Z}) + \widehat{g}(\mathbf{Z})]$ for all $\varepsilon < \varepsilon_0$, $n \leq 6$ and for all \mathbf{A} with $\|\mathbf{A}\| \leq 1$. If $n = 4$ and $\overline{\lambda^1} = \overline{\lambda^3} = \overline{\lambda^5} = 0$, then the first three terms vanish (Lemma 2); the bound (3.9) then follows by replacing \mathbf{A} with $\mathbf{A}/\|\mathbf{A}\|$ and letting $\varepsilon = \|\mathbf{A}\|$. The bound (3.10) can be established in a similar manner.

REFERENCES

- [1] J. ABATE AND W. WHITT, *The Fourier-series method for inverting transforms of probability distributions*, Queueing Syst., 10 (1992), pp. 5–88.
- [2] J. ABATE AND W. WHITT, *Numerical inversion of Laplace transforms of probability distributions*, ORSA J. Comput., 7 (1995), pp. 36–43.
- [3] M. ABRAMOWITZ AND I. A. STEGUN, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, U.S. Government Printing Office, Washington, D.C., 1964.
- [4] L. ANDERSEN, S. BASU, AND J. SIDENIUS, *All your hedges in one basket*, Risk, November (2003), pp. 67–72.
- [5] L. ANDERSEN AND J. SIDENIUS, *Extensions to the Gaussian copula: Random recovery and random factor loadings*, J. Credit Risk, 1 (2004/05), pp. 29–70.
- [6] X. BURTSCHHELL, J. GREGORY, AND J.-P. LAURENT, *A comparative analysis of CDO pricing models under the factor copula framework*, J. Derivatives, 16 (2009), pp. 9–37.
- [7] X. BURTSCHHELL, J. GREGORY, AND J.-P. LAURENT, *Beyond the Gaussian copula: Stochastic and local correlation*, J. Credit Risk, 3 (2007), pp. 31–62.

- [8] P. CARR AND D. MADAN, *Option valuation using the fast Fourier transform*, J. Comput. Finance, 3 (1999), pp. 463–520.
- [9] M. B. GORDY, *Saddlepoint approximation of CreditRisk+*, J. Bank. Finance, 26 (2002), pp. 1335–1353.
- [10] J. GREGORY AND J.-P. LAURENT, *I will survive*, RISK, June (2003), pp. 103–107.
- [11] J. GREGORY AND J.-P. LAURENT, *In the core of correlation*, RISK, October (2004), pp. 87–91.
- [12] H. HAAF, O. REISS, AND J. SCHOENMAKERS, *Numerically stable computation of CreditRisk+*, in *CreditRisk+ in the Banking Industry*, M. Gundlach and F. Lehrbass, eds., Springer-Verlag, Berlin, 2005, pp. 67–76.
- [13] J. HULL AND A. WHITE, *Valuing credit derivatives using an implied copula approach*, J. Derivatives, 14 (2006), pp. 8–28.
- [14] I. ISCOE, K. JACKSON, A. KREININ, AND X. MA, *Pricing Correlation-Dependent Derivatives Based on Exponential Approximations to the Hockey Stick Function*, working paper, Department of Computer Science, University of Toronto, Toronto, ON, Canada, 2007.
- [15] J.-P. LAURENT AND J. GREGORY, *Basket default swaps, CDOs, and factor copulas*, J. Risk, 7 (2005), pp. 103–122.
- [16] D. LI, *On default correlation: A copula function approach*, J. Fixed Income, March (2000), pp. 115–118.
- [17] R. LORD AND C. KAHL, *Optimal Fourier inversion in semi-analytical option pricing*, J. Comput. Finance, 10 (2007), pp. 1–30.
- [18] R. MARTIN, K. THOMPSON, AND C. BROWNE, *Price and probability*, Risk Magazine, January (2001), pp. 115–117.
- [19] R. C. MERTON, *On the pricing of corporate debt: The risk structure of interest rates*, J. Finance, 29 (1974), pp. 449–470.
- [20] L. SCHLOEGL, *Modelling Correlation Skew via Mixing Copulae and Uncertain Loss at Default*, presentation at the Credit Workshop, Isaac Newton Institute, Cambridge, UK, 2005.