

MULTISCALE INTENSITY MODELS AND NAME GROUPING FOR VALUATION OF MULTI-NAME CREDIT DERIVATIVES

EVAN PAPAGEORGIOU AND RONNIE SIRCAR

ABSTRACT. The pricing of collateralized debt obligations and other basket credit derivatives is contingent upon (i) a realistic modeling of the firms' default times and the correlation between them, and (ii) efficient computational methods for computing the portfolio loss distribution from the individual firms' default time distributions. *Factor models*, a widely-used class of pricing models, are computationally tractable despite the large dimension of the pricing problem, thus satisfying issue (ii), but to have any hope of calibrating CDO data, numerically intense versions of these models are required. We revisit the intensity-based modeling setup for basket credit derivatives and, with the aforementioned issues in mind, we propose improvements (a) via incorporating fast mean-reverting stochastic volatility in the default intensity processes, and (b) by considering homogeneous groups within the original set of firms. This can be thought of as a hybrid of top-down and bottom-up approaches. We present a calibration example, including data in the midst of the 2008 financial credit crisis, and discuss the relative performance of the framework.

1. INTRODUCTION

Portfolio credit derivatives have attracted considerable attention over the last decade among institutional investors due to the high demand for products that extract high yields from the default risk of firms. Multi-name credit derivatives, and in particular collateralized debt obligations (CDOs), provide a complex synthesis of risks associated with default events. The difficulty in the valuation of such financial derivatives is not only because of the large number of firms involved and the heterogeneity among their credit qualities, but also from the fact that the firms' default events often seem to exhibit significant and time-varying correlation.

There is already a large literature on CDOs concerning their financial purpose and the mathematical models used to determine their value. We present here an extension to the current methodology for correlating default intensities in an attempt to price CDOs and variants of this type. The pricing of multi-name credit derivatives requires (i) realistic modeling of the firms' default times and the correlation between them, and (ii) efficient computational methods for computing the portfolio loss distribution. *Dynamic factor models*, a widely-used class of pricing models in the literature, can be computationally tractable despite the large

Date: First draft: June 27, 2007; revised 3 March, 2008; current version 7 December, 2008.

JEL Classification. G12, G13.

Key words and phrases. Collateralized debt obligations, intensity-based model, stochastic volatility, asymptotic approximation, multiple time scales, homogeneous-group factor models, bottom-up, top-down.

Work partially supported by NSF grant DMS-0456195.

dimension of the pricing problem, thus satisfying issue (ii), but to have any hope of calibrating CDO data, numerically intense versions of these models have to be implemented. In the approach we analyze here, we start from intensity-based models for default risk, and with the aforementioned issues in mind, we propose improvements (a) via incorporating fast mean-reverting stochastic volatility in the default intensity processes, and (b) by breaking the original set of firms into homogeneous groups.

The most widely-used framework for *single-name* default modeling is the *intensity-based* setup, in which the arrival of default of a firm is described by the first jump of a *Cox process* (or *doubly-stochastic Poisson process*). We refer to Lando (2004) for background and references. Extensions of this *dynamic* framework to *multi-name* credit derivatives were proposed by Duffie and Gârleanu (2001), and more recently by Mortensen (2006) in the class of *common factor models*. We elaborate on this in the next section. Other approaches for pricing basket credit derivatives, such as Markov chain models were studied by Frey and Backhaus (2006), Di Graziano and Rogers (2006), Hurd and Kuznetsov (2005), and Bielecki et al. (2006). Fouque et al. (2008) have proposed a stochastic volatility structural model for CDO valuation.

Let us also note that the aforementioned models are part of the *bottom-up* framework of multi-name credit derivatives pricing models, since we build the loss distribution from the distributions of the default times of the individual firms. On the contrary, *top-down* models describe directly the dynamics of the loss distribution of the portfolio without explicit specification of the constituent single name behaviors. The latter class of models are investigated by Errais et al. (2007) and Arnsdorf and Halperin (2007), among others. The model we propose here can be viewed as intermediate between these approaches.

In what follows, we give a brief guided tour of factor models, introduce our homogeneous group structure and explain the computational benefits.

1.1. Review of Common Factor Models. We begin with a cautionary remark. Due to the sudden emergence of research papers on credit risk, and the simultaneous contribution of multiple authors on a specific subject, there is no set terminology for all pricing frameworks. The very term ‘factor model’ is used by various authors with different definitions, and sometimes it is meant to convey the idea of *mixture* (as in Frey and McNeil, 2003) or other times simply all models with certain independence assumptions (see Andersen and Sidenius, 2005). We adopt the latter and revisit their properties below.

Factor models have been used for at least a decade, with the technical reports of Credit Suisse First Boston (1997) and Gupton et al. (1997), where the notions of *systematic* and *idiosyncratic* credit risks were introduced. Loosely speaking, factor models assume that a firm’s probability of default depends on two components: *firm-specific* (also known as *idiosyncratic*), and *market-wide* (referred to also as *systematic*). Additionally, the conditional default events of any two firms in the portfolio, given the systematic factors, are independent. These two assumptions are important in obtaining computationally tractable pricing expressions, as we see below.

Factor models can be separated into *static models* or *dynamic models*. Static models map the single-name default times to the joint distribution of default times using a copula. The most popular choice is the Gaussian copula introduced by Li (2000), and we refer to Cherubini et al. (2004), Schönbucher (2003) and Hull and White (2004) for details.

On the other hand, dynamic models can provide the portfolio loss distribution process either from the single-name loss processes (bottom-up), without the need for an ad-hoc interdependence relation such as copulas, or by direct specification (top-down). As is usually the case in the bottom-up approach, we choose intensity-based models for the single-name default time process and we decompose the default intensity process into the sum of two other independent processes: idiosyncratic and systematic components. We put this in more concrete terms next.

1.2. Modeling Setup and Examples. Let $(\Omega, \mathcal{H}, \mathbb{P})$ be a probability space, and suppose there are n Cox processes $N^i = (N_t^i)_{t \geq 0}$ with non-negative (and non-trivial) intensity processes $\lambda^i = (\lambda_t^i)_{t \geq 0}$, $i = 1, \dots, n$. Here, \mathbb{P} is a risk-neutral pricing measure consistent with market prices of credit derivatives. We are not interested in the historical measure for our valuation problem. Of course many risk-neutral measures may be consistent with the finite data observed in the market. As is standard, calibration of a stochastic model has the effect of selecting one of these measures. The process N^i is dubbed the *default process* to denote that the first arrival from it signifies the *default* of firm i , for every $i = 1, \dots, n$. We define τ_i to be the *time of default* of firm i , or $\tau_i := \inf\{t > 0 : N_t^i > 0\}$. Then, if we define the filtration $\mathcal{F}^i = (\mathcal{F}_t^i)_{t \geq 0}$ via

$$\mathcal{F}_t^i := \sigma \{ \lambda_s^i; 0 \leq s \leq t \},$$

which describes the history of the default intensity process of the i th firm, the default time τ_i satisfies

$$\mathbb{P}\{\tau_i > t \mid \mathcal{F}_s^i\} = \mathbf{1}_{\{\tau_i > s\}} \mathbb{E} \left[\exp \left\{ - \int_s^t \lambda_u^i du \right\} \mid \mathcal{F}_s^i \right], \quad 0 \leq s \leq t.$$

We define the *loss process* $L = (L_t)_{t \geq 0}$ by

$$L_t := \sum_{i=1}^n \mathbf{1}_{\{\tau_i \leq t\}}, \quad t \geq 0,$$

which counts the number of defaulted firms until time t . The path $t \mapsto L_t$ starts from zero, is right-continuous, and it increases only by jumps of unit size (because in this class of models there cannot be simultaneous defaults a.s.).

As we show in Section 3.1, the pricing of portfolio derivatives amounts to specifying the probability law of the process L , which we refer to as the *loss distribution*. We give some examples below of choices for the intensity processes and the resulting loss distributions to enhance our intuition about the loss process.

Example 1.1 (Deterministic and identical). If every default intensity process is deterministic (but time-dependent) and is equal to each other, i.e., we have $\lambda_t^1 \equiv \dots \equiv \lambda_t^n \equiv \lambda(t)$, the loss

distribution is obviously binomial, with mass function $m \mapsto \mathbb{P}\{L_t = m\}$ given by

$$(1.1) \quad \mathbb{P}\{L_t = m\} = \binom{n}{m} \left(1 - e^{-\int_0^t \lambda(s) ds}\right)^m \left(e^{-\int_0^t \lambda(s) ds}\right)^{n-m}, \quad m = 0, \dots, n.$$

Example 1.2 (Independent and identically distributed). If the default intensity processes are independent and identically distributed, the loss distribution remains binomial, but the mass function now becomes

$$\mathbb{P}\{L_t = m\} = \binom{n}{m} \left(1 - \mathbb{E}\left[e^{-\int_0^t \lambda_s ds}\right]\right)^m \mathbb{E}\left[e^{-\int_0^t \lambda_s ds}\right]^{n-m}, \quad m = 0, \dots, n,$$

where the process $\lambda = (\lambda_t)_{t \geq 0}$ is equal in distribution to the independent and identically distributed intensity processes $\lambda^1, \dots, \lambda^n$.

Example 1.3 (Independent). If the default intensity processes are independent but not necessarily identically distributed, the time- t loss distribution $(P^n(m))_{m=0, \dots, n}$, can be built up by first selecting an arbitrary ordering of the firms. Then, let $P^k(m)$ be the probability of m defaults from the pool of k firms ($k = 1, \dots, n$) which can be calculated recursively via

$$\begin{aligned} P^{k+1}(m) &= P^k(m)(1 - p_{k+1}) + P^k(m-1)p_{k+1}, \quad m = 1, \dots, k+1, \\ P^{k+1}(0) &= P^k(0)(1 - p_{k+1}), \end{aligned}$$

for $k = 1, \dots, n-1$, and

$$\begin{aligned} P^0(m) &= \mathbf{1}_{\{m=0\}}, \quad m = 0, \dots, n, \\ P^k(m) &= 0, \quad m > k, \end{aligned}$$

where p_k is the probability the k th firm defaults by time t , or

$$p_k = 1 - \mathbb{E}\left[e^{-\int_0^t \lambda_s^k ds}\right], \quad k = 1, \dots, n.$$

This is the recursion algorithm suggested by Andersen et al. (2003) and Hull and White (2004) for the computation of the loss distribution. Notice that the distribution $(P^n(m))_{m=0, \dots, n}$ is always the same, regardless of the ordering of firms that is used for this recursion algorithm (but $(P^k(m))_{m=0, \dots, k}$, $k < n$ do not have to be).

Example 1.4 (Identical). If there is a common intensity process driving *all* Cox processes, i.e., there is $\lambda = (\lambda_t)_{t \geq 0}$ with $\lambda_t^1 = \dots = \lambda_t^n = \lambda_t$ a.s., the *conditional loss distribution* given the path of the common intensity $(\lambda_s)_{s \leq t}$ is binomial as argued in (1.1). This implies then that

$$(1.2) \quad \mathbb{P}\{L_t = m\} = \mathbb{E}\left[\binom{n}{m} \left(1 - e^{-\int_0^t \lambda_s ds}\right)^m \left(e^{-\int_0^t \lambda_s ds}\right)^{n-m}\right], \quad m = 0, \dots, n.$$

Of course, the loss distribution is no longer binomial in general.

The assumption of independent default intensities as illustrated in Example 1.2 simplifies the computation of the loss distribution to the computation of a Laplace transform of the integrated intensity process. On the other hand, the situation described in Example 1.4 is simplistic as well, but, as a starting point for a more complicated approach, it is far from the unrealistic binomial case of Example 1.2. The computation of the loss distribution in

the common intensity process of Example 1.4 may also be reduced to the computation of the Laplace transform of the integrated intensity process if we rewrite the probability in (1.2) as

$$\mathbb{P}\{L_t = m\} = \binom{n}{m} \sum_{j=0}^m \binom{m}{j} (-1)^{m-j} \mathbb{E} \left[e^{-(n-j) \int_0^t \lambda_s ds} \right], \quad m = 0, \dots, n.$$

Summations of this type where the terms have alternating signs are known as *Euler-Maclaurin sums*. As n increases, the evaluation of these sums becomes computationally intensive due to round-off numerical errors committed by subtracting large numbers that differ only in trailing decimals. Even though there exist efficient algorithms for accurate numerical computation of these summations, it is necessary to limit the number of firms in the portfolio, n , by 30. (Also, there is a numerical error generated by the multiplication of large binomial coefficients with the small Laplace transform terms.) Obviously, this is a disadvantage considering that the typical size of the name pool is 125 (for the CDX and iTraxx indices) and up to 600 for bespoke CDOs.

Example 1.5 (Common Factor). The default intensity process λ^i of firm i , $i = 1, \dots, n$, is decomposed into an idiosyncratic and a systematic component:

$$\lambda^i = X^i + c_i Z,$$

where X^i and Z are positive and independent processes, and c_1, \dots, c_n are positive numbers. The processes X^1, \dots, X^n are independent, as well. With these assumptions, the market-wide behavior of default risk influences all firms increasingly with its c -coefficient, while there is also a firm-specific component to every firm's default time. The name-to-name correlation is imposed only through the dependence of each λ^i on the market factor Z .

Let us define $U = (U_t)_{t \geq 0}$ to be the *integrated systematic* process, i.e., $U_t = \int_0^t Z_s ds$. When the processes X^1, \dots, X^n have the same law, the loss process is computed with the binomial distribution of Example 1.2, after conditioning on the path of the common process Z . This leads to

$$\mathbb{P}\{L_t = m\} = \mathbb{E} \left[\binom{n}{m} \left(1 - \mathbb{E} \left[e^{-\int_0^t \lambda_s ds} \mid U_t \right] \right)^m \mathbb{E} \left[e^{-\int_0^t \lambda_s ds} \mid U_t \right]^{n-m} \right].$$

When the processes X^1, \dots, X^n do not, necessarily, have the same law (but are still independent), the conditional loss distribution is computed according to Example 1.3. Usual stochastic processes considered for the dynamics of X^1, \dots, X^n and Z are affine diffusions or affine jump-diffusions so that their integrated versions have closed-form, or nearly-closed-form, expressions for the Laplace transforms required. This methodology was originally suggested by Duffie and Gârleanu (2001), and it was later extended by Mortensen (2006) for a parsimonious calibration of the model.

2. HOMOGENEOUS GROUPS AND MULTISCALE STOCHASTIC VOLATILITY DEFAULT INTENSITIES

In this section, we detail our framework which comprises the partitioning of the firms into groups, and allowing for stochastic volatility in the firms' default intensity processes.

2.1. Homogeneous Group Structure. The two modeling situations that are described in Examples 1.2 and 1.4 are intuitively at opposite extremes. We take advantage of their computational efficiencies by breaking the large number of firms into a smaller number of groups and by linking the default intensities across the groups by a common factor.

Assumption 2.1. *There are k homogeneous groups of firms ($k < n$), and each of the n firms belongs in one group only. We denote by n_i the number of firms in the i th group, $i = 1, \dots, k$.*

Assumption 2.2. *Within each group, the firms share a common intensity of default, therefore given a default from the group each (remaining) firm within that group is equally likely to be the defaulted one.*

Let $L^i = (L_t^i)_{t \geq 0}$ be the loss process corresponding to the firms of group i , $i = 1, \dots, k$, and obviously the total loss process $L = L^1 + \dots + L^k$. For $i = 1, \dots, k$, we redefine $\lambda^i = (\lambda_t^i)_{t \geq 0}$ to denote the default intensity process shared by the firms of group i .

Assumption 2.3. *The default intensity process shared by all firms of the homogeneous group i , $\lambda^i = (\lambda_t^i)_{t \geq 0}$ is given by*

$$(2.1) \quad \lambda_t^i = X_t^i + c_i Z_t, \quad i = 1, \dots, k.$$

The idiosyncratic factors X^1, \dots, X^k are independent from each other and independent of the systematic factor Z . The processes X^1, \dots, X^k and Z are non-negative almost surely, and the parameters c_1, \dots, c_k are positive constants.

Assumptions 2.1 and 2.2 were also adopted by Frey and Backhaus (2006) in the context of a mean-field interaction model.

In our setup, the default times of the firms within a given group are heavily correlated since they share identical intensities, and the loss distributions of different groups are correlated through the common factor Z . As we shall see in the calibration section, the presence of the systematic factor enhances enormously the capability of the model to capture the data.

We define the integrated common factor U by

$$U_t = \int_0^t Z_s ds, \quad t > 0.$$

Then, the conditional mass function of the loss in the i th group is given by:

$$(2.2) \quad \mathbb{P}\{L_t^i = m \mid U_t = v\} = \binom{n_i}{m} \sum_{j=0}^m \binom{m}{j} (-1)^{m-j} e^{-c_i(n_i-j)v} \mathbb{E} \left[e^{-(n_i-j) \int_0^t X_s^i ds} \right],$$

for $m = 0, \dots, n_i$. The conditional distribution of the entire portfolio is then given by the convolution

$$(2.3) \quad \mathbb{P}\{L_t = m \mid U_t = v\} = \sum_{i_1=0}^m \dots \sum_{i_{k-1}=0}^{m-i_1-\dots-i_{k-2}} \mathbb{P}\{L_t^1 = i_1 \mid U_t = v\} \dots \\ \cdot \mathbb{P}\{L_t^{k-1} = i_{k-1} \mid U_t = v\} \mathbb{P}\{L_t^k = m - i_1 - \dots - i_{k-1} \mid U_t = v\},$$

where $m = 0, \dots, n$. (Recall also that for every $i = 1, \dots, k$, $\mathbb{P}\{L_t^i = m \mid U_t = v\} = 0$, for $m \notin \{0, \dots, n_i\}$.) The numerical computation of the nested sums is very intensive due to

the large number of firms, n , and number of groups, k . Instead, we compute the conditional distribution of L_t given U_t from the distributions of L_t^1, \dots, L_t^k by inverting the conditional generating function:

$$(2.4) \quad \mathbb{E} [z^{L_t} | U_t = v] = \mathbb{E} [z^{L_t^1} | U_t = v] \cdots \mathbb{E} [z^{L_t^k} | U_t = v].$$

A popular inversion method of the left-hand side of (2.4) is the fast Fourier transform (FFT) method which we use here and explain in Appendix A.2.

From (2.2) we see that we need to compute the Laplace transform of the integrated idiosyncratic component X^i of the intensity process λ^i . The class of affine continuous stochastic processes yield closed-form expressions for such expectations. We discuss such models and important extensions in the next section. Finally, the portfolio loss distribution is computed from

$$(2.5) \quad \mathbb{P}\{L_t = m\} = \int_0^{+\infty} \mathbb{P}\{L_t = m | U_t = v\} \mathbb{P}\{U_t \in dv\}, \quad m = 0, \dots, n.$$

In Section 2.2, we describe our stochastic volatility default intensity process, and an asymptotic expansion, along the lines of Cotton et al. (2004) for interest rate derivatives, that leads to analytic approximations for the group loss distributions. CDOs have a complicated structure which we describe in Section 3.1, and we offer a calibration example from standardized index data in Section 3.2. The final Section 4 concludes and discusses areas for further research.

2.2. Stochastic Volatility in the Default Intensity Process. Volatility models are common in the fixed income literature, paralleling their equity counterparts. Stochastic volatility extensions on LIBOR and other interest-rate models have been suggested for pricing, and we refer to Musiela and Rutkowski (2007) for details. A different class of stochastic volatility models offer asymptotic approximations for the pricing functions for stochastic volatility driven by fast mean-reverting processes. Such models were studied in Fouque et al. (2000) for equity derivatives, and Cotton et al. (2004) for interest-rate derivatives. Here, we extend the latter framework of stochastic volatility models to the default intensity process.

Modeling the default intensity processes in a Cox process setup with a single-scale diffusion process was found to be inadequate in producing appropriate loss distributions to capture real data (see Duffie and Gârleanu, 2001; Mortensen, 2006). In particular, the market-implied loss distribution exhibits longer tails than the ones corresponding to diffusion processes. Adding jumps yields satisfactory results. As we argue below, the introduction of stochastic volatility in the asymptotic approximation is enough to allow for a *heavier tail* in the loss distribution, which offsets the need for jump characteristics, and maintains closed-form expressions (up to multiscale asymptotic approximations) for the conditional loss distribution.

In a previous work (Papageorgiou and Sircar, 2008) we presented related asymptotic approximations for the prices of single name credit derivatives such as defaultable bonds, options on defaultable bonds, and credit default swaps using multiscale intensity models.

2.2.1. *Fast Mean-Reverting Stochastic Volatility.* Let us fix a group i ($i = 1, \dots, k$). The model for the two components of the default intensity process in (2.1) is described by the system of stochastic differential equations:

$$(2.6) \quad dX_t^i = \alpha_i(\bar{x}_i - X_t^i) dt + f_i(Y_t^i) \sqrt{X_t^i} dW_t^i,$$

$$(2.7) \quad dY_t^i = \frac{1}{\varepsilon} X_t^i (\bar{y}_i - Y_t^i) dt + \frac{\nu_i \sqrt{2}}{\sqrt{\varepsilon}} \sqrt{X_t^i} dW_t^{Y^i},$$

$$dZ_t = \alpha_z(\bar{z} - Z_t) dt + \sigma_z \sqrt{Z_t} dW_t^Z.$$

Here, W^i and W^{Y^i} are Wiener processes such that

$$d\langle W^i, W^{Y^i} \rangle_t = \rho_i dt, \quad d\langle W^i, W^{Y^j} \rangle_t = 0, \quad d\langle W^i, W^j \rangle_t = 0, \quad d\langle W^{Y^i}, W^{Y^j} \rangle_t = 0,$$

for $i = 1, \dots, k, j = 1, \dots, k, i \neq j$, and ρ_i in $[-1, 1]$. Due to the factor model Assumption 2.3 the Wiener process W^Z is independent of W^i and W^{Y^i} .

The idiosyncratic intensity process X^i is a square-root diffusion process with stochastic volatility; it is a mean-reverting process with rate of mean reversion α_i , and level of mean reversion \bar{x}_i . The diffusion coefficient is driven by another mean-reverting process Y^i via a function f_i which we assume to be positive and bounded in order to (i) have strong solutions to (2.6), and (ii) to satisfy certain growth conditions needed for the proof of the approximation theorem later. The condition $2\alpha_i \bar{x}_i > f_i(\cdot)^2$ is sufficient for X^i to remain strictly positive, and the justification follows the one from the constant volatility case (see Feller, 1951).

Notice the presence of the small positive number ε in (2.7) which makes Y^i a *fast mean-reverting process* (because then $1/\varepsilon$ is large). Furthermore, we choose to correlate X^i and Y^i not only through their Wiener processes, but also from the presence of X^i in (2.7). We do so in this specific way to construct closed form approximations for the Laplace transforms of the integrated X^i process. The systematic component, Z , is modeled as a square-root diffusion with constant volatility σ_z . It is possible to allow Z to have stochastic volatility also, to thicken the tails of the mixing distribution of the systemic factor. If this volatility process is fast mean-reverting, a closed-form asymptotic approximation is readily computable. This extension adds further calibration parameters to the many we already have, and we do not pursue it here.

2.2.2. *Asymptotic Approximation.* We start with the Laplace transform of $\int_t^T X_s^i ds$ that is used to compute the conditional (on the common factor Z) group loss distribution and then the total loss portfolio loss distribution in Section 2.3. In what follows we simplify the notation for the processes involved and instead of X^i and Y^i we write X and Y , respectively. The stochastic differential equations (2.6) and (2.7) remain the same with the appropriate drop of the index i .

We construct an approximation to the type of expectation in expression (2.2):

$$(2.8) \quad u(t, T, x, y; m) := \mathbb{E} \left[e^{-m \int_t^T X_s ds} \mid X_t = x, Y_t = y \right], \quad t \leq T.$$

The function u satisfies the partial differential equation (PDE)

$$(2.9) \quad \begin{aligned} \mathcal{L}u(t, T, x, y; m) &= 0, \quad t < T, \quad x \in \mathbb{R}_+, \quad y \in \mathbb{R}, \\ u(T, T, x, y; m) &= 1, \end{aligned}$$

with

$$\begin{aligned} \mathcal{L} &:= \frac{1}{\varepsilon} \mathcal{L}_0 + \frac{1}{\sqrt{\varepsilon}} \mathcal{L}_1 + \mathcal{L}_2, \\ \mathcal{L}_0 &:= x \left(\nu^2 \frac{\partial^2}{\partial y^2} + (\bar{y} - y) \frac{\partial}{\partial y} \right), \\ \mathcal{L}_1 &:= \sqrt{2} \rho \nu f(y) x \frac{\partial^2}{\partial x \partial y}, \\ \mathcal{L}_2 &:= \frac{\partial}{\partial t} + \frac{1}{2} f(y)^2 x \frac{\partial^2}{\partial x^2} + \alpha(\bar{x} - x) \frac{\partial}{\partial x} - mx. \end{aligned}$$

Recall that in the absence of stochastic volatility, that is, $f(y) \equiv \sigma > 0$, the expression for u in (2.8) is given by

$$(2.10) \quad u(t, T, x, y; m) = u(t, T, x; m) = A(T - t) e^{-B(T-t)x},$$

where A and B satisfy the ordinary differential equations (ODEs):

$$\begin{aligned} \frac{A'}{A} + \alpha \bar{x} B &= 0, \\ B' + \frac{1}{2} \sigma^2 B^2 + \alpha B - m &= 0. \end{aligned}$$

Using the initial condition $A(0) = 1$ and $B(0) = 0$ we have the solutions to these ODEs:

$$(2.11) \quad A(s) = \left[\frac{2\gamma e^{(\alpha+\gamma)s/2}}{(\alpha+\gamma)(e^{\gamma s} - 1) + 2\gamma} \right]^{\frac{2\alpha\bar{x}}{\sigma^2}},$$

$$(2.12) \quad B(s) = \frac{2m(e^{\gamma s} - 1)}{(\alpha+\gamma)(e^{\gamma s} - 1) + 2\gamma},$$

$$(2.13) \quad \gamma = \sqrt{\alpha^2 + 2m\sigma^2}.$$

For more on the affine characterization of SDEs and their resulting Laplace transforms, we refer to Brigo and Mercurio (2006).

Initially, we are interested in an expansion of the form

$$(2.14) \quad u = u_0 + \sqrt{\varepsilon} u_1 + \varepsilon u_2 + \varepsilon^{3/2} u_3 + \dots,$$

where each term u_0, u_1, \dots is independent of ε . We will describe the sense of convergence as $\varepsilon \downarrow 0$ in Theorem 2.5.

We plug (2.14) into the left-hand side of (2.9) and since $x > 0$, the $1/\varepsilon$ terms yield the ODE $\mathcal{L}_0 u_0 = 0$. We choose u_0 not to depend on y , i.e., $u_0(t, T, x, y; m) = u_0(t, T, x; m)$, since all other solutions grow unreasonably as $|y| \rightarrow \infty$ (see Fouque et al., 2003). Similarly, the $1/\sqrt{\varepsilon}$ terms of the expansion of (2.9) yield a similar ODE: $\mathcal{L}_0 u_1 = 0$, after recognizing that \mathcal{L}_1 takes a derivative in y , which makes $\mathcal{L}_1 u_0 = 0$ because u_0 does not depend on y . Therefore, we choose u_1 not to depend on y , or $u_1(t, T, x, y; m) = u_1(t, T, x; m)$.

The ε^0 terms of the expansion of (2.9) yield the ODE

$$x \mathcal{L}_0 u_2 + \mathcal{L}_1 u_1 + \mathcal{L}_2 u_0 = 0,$$

and since u_1 does not depend on y it simplifies to

$$(2.15) \quad \mathcal{L}_0 u_2 = -\frac{1}{x} \mathcal{L}_2 u_0, \quad t < T, \quad x > 0.$$

This ODE is a Poisson equation in u_2 with respect to y , and in order for it to have a solution the right-hand side of (2.15) needs to satisfy the *Fredholm alternative* (or *centering condition*):

$$(2.16) \quad \left\langle \frac{1}{x} \mathcal{L}_2 u_0 \right\rangle = 0, \quad t < T.$$

Here the angle brackets $\langle \cdot \rangle$ denote integration with respect to the invariant distribution of the process that has \mathcal{L}_0 as its infinitesimal generator—namely an Ornstein-Uhlenbeck process, which has the Gaussian distribution with mean \bar{y} and variance ν^2 .

Since x is strictly positive and u_0 was chosen not to depend on y the centering condition (2.16) leads to the PDE problem

$$(2.17) \quad \langle \mathcal{L}_2 \rangle u_0 = 0, \quad t < T,$$

where

$$\begin{aligned} \langle \mathcal{L}_2 \rangle &:= \frac{\partial}{\partial t} + \frac{1}{2} \overline{\sigma^2} x \frac{\partial^2}{\partial x^2} + \alpha(\bar{x} - x) \frac{\partial}{\partial x} - mx, \\ \overline{\sigma^2} &:= \langle f(y)^2 \rangle = \int_{\mathbb{R}} f(y)^2 \frac{1}{\sqrt{2\pi\nu^2}} \exp\left\{-\frac{(y-\bar{y})^2}{2\nu^2}\right\} dy. \end{aligned}$$

As we noted in (2.10) this PDE paired with the terminal condition $u_0(T, T, x; m) = 1$, has solution

$$(2.18) \quad u_0(t, T, x; m) = A(T-t)e^{-B(T-t)x}, \quad t \leq T,$$

where A and B were defined in (2.11) and (2.12) after replacing the term σ^2 with $\overline{\sigma^2}$.

For the expression of the u_1 term, we rewrite (2.15) as

$$(2.19) \quad \mathcal{L}_0 u_2 = -\frac{1}{x} (\mathcal{L}_2 - \langle \mathcal{L}_2 \rangle) u_0.$$

Then we set φ to be a solution of the ODE

$$(2.20) \quad \mathcal{L}_0 \varphi(y) = f(y)^2 - \langle f(y)^2 \rangle.$$

The $\sqrt{\varepsilon}$ terms of the expanded (2.9) yield the Poisson equation in y for u_3

$$x \mathcal{L}_0 u_3 + \mathcal{L}_1 u_2 + \mathcal{L}_2 u_1 = 0,$$

the centering condition of which is

$$(2.21) \quad \langle \mathcal{L}_1 u_2 + \mathcal{L}_2 u_1 \rangle = 0,$$

because $x > 0$. Since we chose u_1 to be independent of y , and by using the ODE (2.19) for u_2 we rewrite the centering condition as

$$(2.22) \quad \langle \mathcal{L}_2 \rangle u_1 = \mathcal{A} u_0,$$

where the operator \mathcal{A} is given by

$$\mathcal{A} := \left\langle \mathcal{L}_1 \mathcal{L}_0^{-1} \frac{1}{x} (\mathcal{L}_2 - \langle \mathcal{L}_2 \rangle) \right\rangle.$$

Finally, because of the definition of φ in (2.20), the right-hand side of (2.22) becomes

$$\begin{aligned} \mathcal{A}u_0 &= \left\langle \mathcal{L}_1 \left(\frac{1}{2} \varphi(y) \frac{\partial^2}{\partial x^2} \right) \right\rangle u_0 \\ &= \left\langle \sqrt{2} \rho \nu f(y) x \frac{\partial^2}{\partial x \partial y} \frac{1}{2} \varphi(y) \frac{\partial^2}{\partial x^2} \right\rangle u_0 \\ &= \frac{1}{\sqrt{2}} \rho \nu \langle f \varphi' \rangle x \frac{\partial^3}{\partial x^3} u_0. \end{aligned}$$

By introducing the notation $\tilde{u}_1 := \sqrt{\varepsilon} u_1$, letting $V_1 := \sqrt{\varepsilon} \rho \nu \langle f \varphi' \rangle / \sqrt{2}$, and observing from (2.18) that $\partial^3 u_0 / \partial x^3 = -B^3 u_0$, we write the following PDE problem for \tilde{u}_1 stemming from the PDE (2.22)

$$(2.23) \quad \begin{aligned} \langle \mathcal{L}_2 \rangle \tilde{u}_1(t, T, x; m) &= -V_1 B(T-t)^3 x u_0(t, T, x; m), \quad t < T, \\ \tilde{u}_1(T, T, x; m) &= 0. \end{aligned}$$

To solve the PDE (2.23) we take advantage of the fact that u_0 satisfies the similar problem (2.17). Hence, we make the ansatz

$$(2.24) \quad \tilde{u}_1(t, T, x; m) = (D_1(T-t)x + D_2(T-t))u_0(t, T, x; m),$$

and we get the ODEs for D_1 and D_2 by separating in (2.23) the x -terms and the terms independent of x :

$$(2.25) \quad D_1' + \left(\overline{\sigma^2} B + \alpha \right) D_1 - V_1 B^3 = 0, \quad D_1(0) = 0,$$

$$(2.26) \quad D_2' - \alpha \bar{x} D_2 = 0, \quad D_2(0) = 0.$$

By setting $\theta := \overline{\sigma^2} / (\alpha \bar{x})$, the solution of D_1 in (2.25) is

$$(2.27) \quad D_1(s) = \frac{m^2 V_1}{\gamma^2 \overline{\sigma^2}} e^{-\alpha s} A(s)^\theta \left[\frac{2\gamma^2}{m\alpha \bar{x}} \log(A(s)) + \frac{\alpha + \gamma - 2\gamma e^{\gamma s}}{\gamma e^{\gamma s}} + \frac{(\gamma - \alpha)e^{\gamma s}}{\gamma} + 2\alpha s \right],$$

where A and γ were defined in (2.11) and (2.13) (replacing σ^2 by $\overline{\sigma^2}$). The solution of (2.26) for D_2 is given by

$$(2.28) \quad \begin{aligned} D_2(s) &= \frac{2V_1}{(\overline{\sigma^2})^2} \left[\frac{8\alpha\gamma}{h(s)} + \frac{(\alpha + \gamma)(\gamma + 3\alpha)B(s)}{2m} \right] \log(A(s)) \\ &\quad + \frac{2mV_1\alpha\bar{x}}{\overline{\sigma^2}\gamma(\gamma - \alpha)} \left[\frac{4m(3\gamma^2 + \alpha^2)}{(\alpha + \gamma)h(s)} + (3\gamma + \alpha)B(s) \right] s - \frac{2V_1\alpha\bar{x}(\alpha^2 + 3\gamma^2)B(s)}{(\overline{\sigma^2})^2 \gamma^2}, \end{aligned}$$

where B was defined in (2.12) and

$$h(s) := (\alpha + \gamma)(e^{\gamma s} - 1) + 2\gamma.$$

The approximation of u as defined in (2.8) up to terms of order ε is denoted by u^ε and is given by

$$(2.29) \quad \begin{aligned} u^\varepsilon(t, T, x; m) &= u_0(t, T, x; m) + \tilde{u}_1(t, T, x; m) \\ &= [1 + D_1(T-t)x + D_2(T-t)]A(T-t)e^{-B(T-t)x}, \end{aligned}$$

with A and B defined in (2.11) and (2.12) with $\overline{\sigma^2}$ replacing σ^2 .

Remark 2.4. Notice that the asymptotic approximation for u up to terms of order ε does not depend on the current level of the fast mean-reverting process Y , i.e., on $Y_t = y$. Instead, the parameters of the process Y appear only in the definition of the parameter V_1 , which we will call an *aggregate parameter*.

2.2.3. Accuracy of the Approximation. Regarding the order of approximation, we state the asymptotic behavior of u^ε as $\varepsilon \rightarrow 0$ in the following theorem.

Theorem 2.5. *For fixed $0 < t < T$, fixed $x, y \in \mathbb{R}$, $m = 1, 2, \dots$, and for every $\varepsilon \leq 1$, there exists a positive constant $c < \infty$ that depends on (t, x, y) but not on ε such that*

$$|u(t, T, x, y; m) - u^\varepsilon(t, T, x; m)| \leq c\varepsilon,$$

where u is defined in (2.8) and u^ε in (2.29).

Proof. Recall the defining problem for the u_2 term in (2.19), and let us define the higher order term u_3 of the asymptotic expansion (2.14) via the PDE problem

$$(2.30) \quad \mathcal{L}_0 u_3 = -\frac{1}{x} ((\mathcal{L}_1 u_2 - \langle \mathcal{L}_1 u_2 \rangle) + (\mathcal{L}_2 - \langle \mathcal{L}_2 \rangle) u_1).$$

This PDE is derived from the order $\sqrt{\varepsilon}$ terms of the expansion of the PDE (2.9) by (2.14) and the centering condition (2.21).

Let us define now the higher order approximation to u :

$$\widehat{u}^\varepsilon(t, T, x, y; m) := u^\varepsilon(t, T, x; m) + \varepsilon u_2(t, T, x, y; m) + \varepsilon^{3/2} u_3(t, T, x, y; m),$$

and we form the error term

$$Q(t, T, x, y) := \widehat{u}^\varepsilon(t, T, x, y; m) - u(t, T, x, y; m).$$

Notice that Q satisfies the PDE

$$(2.31) \quad \mathcal{L}Q = \mathcal{L}\widehat{u}^\varepsilon - \mathcal{L}u = \varepsilon Q_1,$$

where

$$Q_1(t, T, x, y) = \mathcal{L}_2 u_2(t, T, x, y; m) + \mathcal{L}_1 u_3(t, T, x, y; m) + \sqrt{\varepsilon} \mathcal{L}_2 u_3(t, T, x, y; m),$$

since u satisfies the defining PDE problem (2.9), and u_0 , u_1 , u_2 , and u_3 are defined by the equations (2.17), (2.22), (2.19), and (2.30) respectively. Observe furthermore that

$$\begin{aligned} Q(T, T, x, y) &= \widehat{u}^\varepsilon(T, T, x; m) - u(T, T, x, y; m) \\ &= \varepsilon (u_2(T, T, x, y; m) + \sqrt{\varepsilon} u_3(T, T, x, y; m)) \\ &=: \varepsilon G(T, x, y). \end{aligned}$$

Then the probabilistic equivalent of the error term Q that satisfies (2.31) is

$$Q(t, T, x, y) = \varepsilon \mathbb{E}_{t,x,y} \left[e^{-\int_t^T X_s ds} G(T, X_T, Y_T) - \int_t^T e^{-\int_t^u X_s ds} Q_1(u, T, X_u, Y_u) du \right],$$

where $\mathbb{E}_{t,x,y}[\cdot] = \mathbb{E}[\cdot \mid X_t = x, Y_t = y]$. Since the function f was assumed to be positive and bounded we can choose φ in (2.20) to be growing at most linearly in $|y|$. This is an important property of solutions of Poisson equations, and we refer to Lemma 4.3 of Fouque et al. (2003) for a detailed proof.

The process X remains strictly positive at all times under our assumptions on f_i, α_i, \bar{x}_i , and so the discounting term $\exp\{-\int_t^T X_s ds\}$ is always between zero and one. We showed in (2.24) that u_1 depends on x not only through its discounting term $\exp\{-Bx\}$ but also through a simple multiplicative relation. It can be easily shown (a matter of algebraic computations in the defining PDEs (2.19) and (2.30)) that the higher order terms u_2 and u_3 also depend on x through the discounting term $\exp\{-Bx\}$ and a polynomial fashion, that is, they can be expressed as

$$\begin{aligned} u_2 &= \left(\hat{D}_1 x^2 + \hat{D}_2 x + \hat{D}_3 \right) u_0, \\ u_3 &= \left(\check{D}_1 x^3 + \check{D}_2 x^2 + \check{D}_3 x + \check{D}_4 \right) u_0, \end{aligned}$$

for some functions $\hat{D}_1, \hat{D}_2, \hat{D}_3, \check{D}_1, \check{D}_2, \check{D}_3, \check{D}_4$ that do not depend on x . Therefore, u_0, u_1, u_2 , and u_3 can be bounded by smooth functions of x .

Notice now that the right-hand side expressions of the defining Poisson equations for the higher order terms u_2 and u_3 in (2.19) and (2.30) grow at most linearly in $|y|$ because either of their dependence on φ , or the fact that they are ODEs in y . Therefore, we can choose u_2 and u_3 as well to grow at most linearly in $|y|$. This shows that Q_1 and G grow at most linearly in $|y|$, are bounded by smooth functions of x , and that Q_1 is bounded by a smooth function of time t as an ODE in y . (The term G does not depend on t .)

Then we have by the triangle inequality (and suppressing the notation at places)

$$\begin{aligned} |u(t, T, x, y; m) - u^\varepsilon(t, T, x; m)| &\leq |u - \hat{u}^\varepsilon| + |\hat{u}^\varepsilon - u^\varepsilon| \\ &= |Q| + \varepsilon |u_2 + \sqrt{\varepsilon} u_3| \\ &\leq (c_1 + c_2) \varepsilon \end{aligned}$$

for some constants c_1, c_2 , which depend on the parameters, but not on ε . This completes the proof. \square

2.2.4. Positivity Preserving Asymptotic Approximation. The asymptotic approximation u^ε in (2.29) may become negative for certain values of the aggregate parameter and large maturity T . Of course, this is normal due to the type of the approximation, but it is a disadvantage nonetheless. To account for this problem, we follow Fouque and Zhou (2006), and we instead assume a multiplicative asymptotic perturbation for u in (2.8) of the form

$$\bar{u}^\varepsilon(t, T, x, y; m) = \hat{u}_0(t, T, x, y; m) [1 + \tanh(\sqrt{\varepsilon} \hat{u}_1(t, T, x, y; m))].$$

The terms \hat{u}_0 and \hat{u}_1 are defined so that

$$u^\varepsilon(t, T, x; m) = \bar{u}^\varepsilon(t, T, x, y; m),$$

i.e., $u_0 + \sqrt{\varepsilon} u_1 = \hat{u}_0(1 + \sqrt{\varepsilon} \hat{u}_1)$, which yields the new approximation for u :

$$(2.32) \quad \begin{aligned} \bar{u}^\varepsilon(t, T, x; m) &= u_0(1 + \tanh(\sqrt{\varepsilon} u_1/u_0)) \\ &= u_0[1 + \tanh(D_1(T-t)x + D_2(T-t))]. \end{aligned}$$

This expression remains positive for all values of V_1 and every time $t \leq T$, by construction. Note also that the functions D_1 and D_2 as defined in (2.27) and (2.28) are directly proportional to V_1 .

The approximation of u by \bar{u}^ε is of the same order as by u^ε as given in Theorem 2.5. Notice that for small x we have $\tanh(x) \approx x - x^3/3$ which then gives for fixed t, T, x, m the following approximation

$$\begin{aligned} \bar{u}^\varepsilon &\approx u_0 \left(1 + \sqrt{\varepsilon} \frac{u_1}{u_0} - \frac{\varepsilon^{3/2}}{3} \left(\frac{u_1}{u_0} \right)^3 \right) \\ &= u^\varepsilon - \varepsilon^{3/2} \frac{u_1^3}{3u_0^2}. \end{aligned}$$

Finally, we have

$$\begin{aligned} |u - \bar{u}^\varepsilon| &\leq |u - u^\varepsilon| + |u^\varepsilon - \bar{u}^\varepsilon| \\ &\leq c_1 \varepsilon - c_2 \varepsilon^{3/2} \\ &= \mathcal{O}(\varepsilon), \end{aligned}$$

where we used Theorem 2.5.

2.3. Summary and Effects of Stochastic Volatility on Loss Distributions. Using the asymptotic approximation \bar{u}^ε as given in (2.32) for u (defined in (2.8)) we provide an approximation for the conditional probability loss density of a group via (2.2), or

$$(2.33) \quad \begin{aligned} \mathbb{P}\{L_t^i = m \mid U_t = v\} &= \binom{n_i}{m} \sum_{j=0}^m \binom{m}{j} (-1)^{m-j} e^{-c_i(n_i-j)v} u^i(0, t, x_i; n_i - j) \\ &\approx \binom{n_i}{m} \sum_{j=0}^m \binom{m}{j} (-1)^{m-j} e^{-c_i(n_i-j)v} [1 + \tanh(D_1^i(t)x_i + D_2^i(t))] u_0^i(0, t, x_i; n_i - j). \end{aligned}$$

The functions D_1^i and D_2^i are defined in (2.27) and (2.28) respectively, modified appropriately to reflect the use of the parameters from the processes X^i and Y^i (in place of the generic parameters of X and Y). Note also that $X_0^i = x_i$.

The total portfolio conditional loss distribution is obtained from the numerical inversion of the conditional generating function (2.4), while the distribution of U_t is taken from the numerical inversion of its characteristic function. Both numerical inversions are done using the fast Fourier transform algorithm. Even though the generating function inversion procedure seems more complicated than the convolution computation (2.3), it is dramatically faster, on the order of 50 times for five equally-sized homogeneous groups of 25 firms per group. See

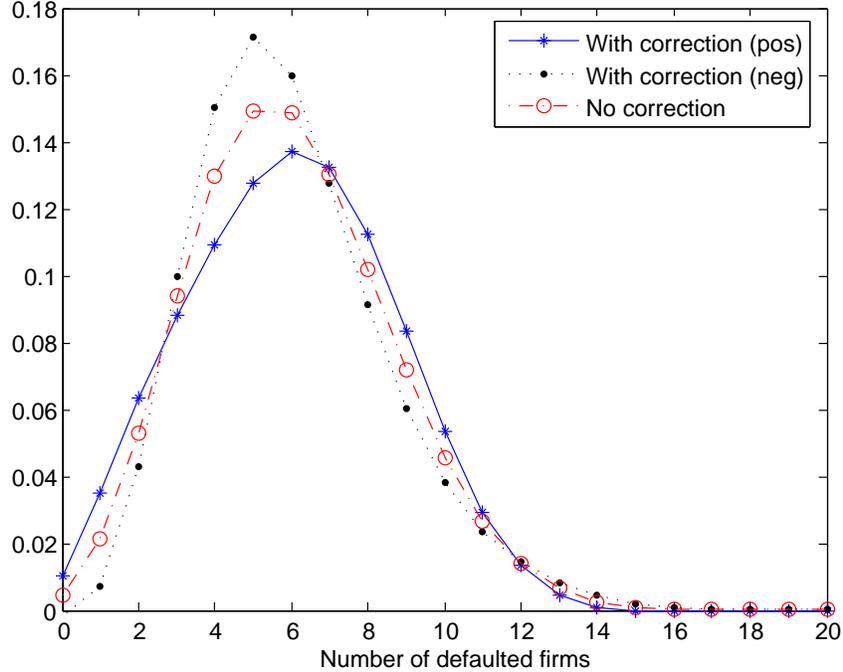


FIGURE 2.1. The probability loss density function within a single homogeneous group $m \mapsto \mathbb{P}\{L_t^i = m\}$ with and without the correction due to stochastic volatility, as shown in (2.5). The solid line corresponds to the density function with positive parameter $V_{1,i} = 0.00015$, while the dotted line has negative $V_{1,i}$ parameter ($V_{1,i} = -0.00015$). They can be compared to the dashed line which corresponds to the constant-volatility density function (i.e., $V_{1,i} = 0$). The other parameters used are $n_i = 20$, $\alpha_i = 0.15$, $\bar{\sigma}_i^2 = 0.1^2$, $\bar{x}_i = 0.015$, $x_i = 0.04$, $\alpha_z = 0.1$, $\bar{z} = 0.005$, $z = 0.004$, $\sigma_z = 0.01$, $t = 5$ years.

Appendix A for more. Finally, the loss distribution is obtained from the numerical evaluation of the integral (2.5).

Remark 2.6. We remark that the recursion algorithm proposed by Andersen et al. (2003) and Hull and White (2004) (illustrated in Example 1.3) is not applicable here because the grouping does not simply reduce the problem to the sum of conditionally independent Bernoulli random variables. Therefore, we need to use inversion of the generating function to obtain the conditional loss distribution of the sum of the random variables L_t^1, \dots, L_t^k , given the integrated systematic factor U_t .

Figure 2.1 exhibits the effect of the stochastic volatility approximation on the probability loss distribution of a group of $n_i = 20$ firms with identical default intensity processes X^i as given in (2.6) and stochastic volatility process $(f_i(Y_t^i))_{t \geq 0}$ and systematic intensity process Z . The parameters used are listed in the caption of the figure. Notice that the stochastic volatility correction allows for the mass of the distribution to be shifted to the right, which is vital for the calibration of the model to market data.

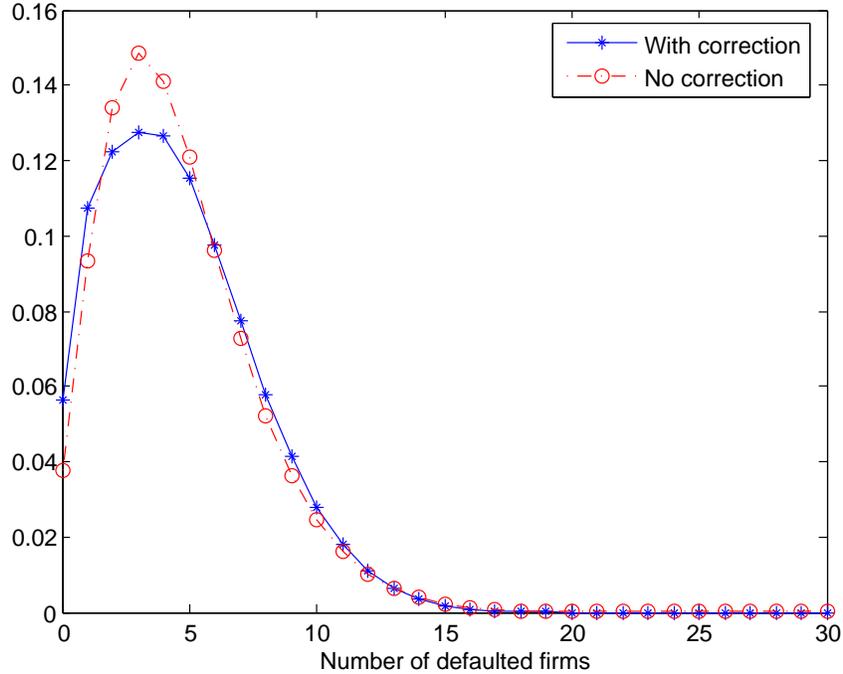


FIGURE 2.2. *The effect of stochastic volatility on the mass function $m \mapsto \mathbb{P}\{L_t = m\}$ of the number of defaulted firms at time t for a basket of 125 firms separated in five equal-size groups of 25 firms per group. The parameters used are $\alpha_1 = \alpha_2 = \alpha_3 = 0.15$, $\alpha_4 = \alpha_5 = 0.1$, $\bar{\sigma}_1^2 = \dots = \bar{\sigma}_5^2 = 0.1^2$, $\bar{x}_1 = \dots = \bar{x}_5 = 0.015$, $x_1 = \dots = x_5 = 0.005$, $V_{1,1} = -0.0001$, $V_{1,2} = -0.00005$, $V_{1,3} = 0.0001$, $V_{1,4} = 0.00015$, $V_{1,5} = 0.0002$, $\alpha_z = 0.1$, $\bar{z} = 0.005$, $z = 0.004$, $\sigma_z = 0.01$, $t = 5$ years. The dashed line illustrates the absence of stochastic volatility approximation by setting $V_{1,1} = \dots = V_{1,5} = 0$. (We graph the mass function over $\{0, 1, \dots, 30\}$ instead of $\{0, 1, \dots, 125\}$.)*

The presence of stochastic volatility in our approximation is also present in the loss process for the entire portfolio as shown in Figure 2.2 which shows the probability mass function $m \mapsto \mathbb{P}\{L_t = m\}$ from the formula (2.5) with five equal-sized groups ($k = 5$) of 25 firms ($n_1 = \dots = n_5 = 25$) for a total of 125. Notice again the mass-shift in the density function compared to the constant volatility case (no correction). The presence of five independent groups results into a lower peak and thicker tail of the density function. This feature allows the model to fit real data very adequately as we will see in the fitting results of Section 3.2.

3. CDOs AND CALIBRATION TO CDX DATA

3.1. CDO Mechanics. We are interested in pricing the *tranches* of a CDO. Let K_1 and K_2 be the *attachment* and *detachment* points of a tranche, or in other words, the endpoints of the range of percentage losses the contract is providing insurance for. The standardized indices CDX and iTraxx have five benchmark tranches, the attachment and detachment points of which are given in Table 3.1. The first tranche is referred to as *equity*, since it is the riskiest and requires the highest premium. The last tranche is called *senior* and in almost all cases it is of very high credit quality. Finally, the three intermediate tranches are called *mezzanine*

tranches. Obviously, the insurance premia are decreasing across the tranches, with the equity tranche having the largest premium, and the senior tranche the smallest.

TABLE 3.1. *The attachment (K_1) and detachment (K_2) points for the tranches of the Dow Jones CDX and iTraxx indices.*

	Equity		Mezz 1		Mezz 2		Mezz 3		Senior	
	K_1	K_2								
CDX	0%	3%	3%	7%	7%	10%	10%	15%	15%	30%
iTraxx	0%	3%	3%	6%	6%	9%	9%	12%	12%	22%

The tranche loss function $\ell_{[K_1, K_2]}$ with attachment and detachment points K_1 and K_2 is defined as

$$\ell_{[K_1, K_2]}(x) = (x - K_1)^+ - (x - K_2)^+, \quad 0 \leq x \leq 1.$$

We assume that each firm has the same notional amount in the portfolio, which without loss of generality we take to be $\$1/n$. We also assume a deterministic recovery $\delta \in [0, 1]$, that is, the loss to the portfolio from each default is $\$(1 - \delta)/n$. Then the portfolio loss at time t is given by

$$\Lambda_t = \frac{(1 - \delta)L_t}{n}, \quad t \geq 0.$$

There are two counterparties in a CDO agreement. The *protection buyer* (usually an investment bank or the originator of the original pool of securitized assets) of the $[K_1, K_2]$ tranche enters into the contract and purchases insurance against any losses between the fractions K_1 and K_2 of the total notional of the portfolio over the time period $[0, T]$. The protection buyer makes M periodic payments at times $t_1 < t_2 < \dots < t_M = T$, which for the case of CDX and iTraxx are quarterly, i.e., $t_1 = 0.25, t_2 = 0.5$, and so on. The *protection seller* (usually institutional investors) receives the payments, and in return will compensate the protection buyer as agreed. Hence, the *pricing* of the tranches of a CDO is the specification of the constant coupon rate the protection buyer will pay to the protection seller on every payment due date. This coupon for the mezzanine and senior tranches (but not equity) is given by (see Lando, 2004, for the CDS case which is easily adaptable to CDO spreads)

$$s = \frac{\sum_{j=1}^M e^{-r(t_j + t_{j-1})/2} (\mathbb{E}[\ell_{[K_1, K_2]}(\Lambda_{t_j})] - \mathbb{E}[\ell_{[K_1, K_2]}(\Lambda_{t_{j-1}})])}{\sum_{j=1}^M (t_j - t_{j-1}) e^{-rt_j} [K_2 - K_1 - \frac{1}{2} (\mathbb{E}[\ell_{[K_1, K_2]}(\Lambda_{t_j})] + \mathbb{E}[\ell_{[K_1, K_2]}(\Lambda_{t_{j-1}})])]}.$$

For the equity tranche it is standard to assume that there is an *up-front* fee in the inception of the contract paid to the protection seller by the protection buyer plus a fixed running premium s of 500 basis points (bps) or 5% of the tranche notional $K_2 - K_1 = K_2$. Therefore

if we define the up-front fee as the fraction s_{eq} of K_2 , we have

$$s_{\text{eq}} = \frac{1}{K_2} \left[\sum_{j=1}^M e^{-r(t_j+t_{j-1})/2} (\mathbb{E}[\ell_{[0,K_2]}(\Lambda_{t_j})] - \mathbb{E}[\ell_{[0,K_2]}(\Lambda_{t_{j-1}})]) - 0.05 \sum_{j=1}^M (t_j - t_{j-1}) e^{-rt_j} \left[K_2 - \frac{1}{2} (\mathbb{E}[\ell_{[0,K_2]}(\Lambda_{t_j})] + \mathbb{E}[\ell_{[0,K_2]}(\Lambda_{t_{j-1}})]) \right] \right].$$

3.2. CDX Data. To exhibit the implementation of the combination of stochastic volatility default intensity models with the homogeneous group structure, we present the results of a preliminary fit to the five-year CDX index tranche spreads.

The CDX is a credit default swap (CDS) index introduced in 2003, and is composed of 125 equally-weighted North American corporations of investment grade credit rating. The firm-components are chosen from seven financial market sectors, and the index is interpreted as the average insurance premium required for protection against the default of an investment-grade company within five years. (There are also seven- and ten-year versions of the index but the five-year CDX quotes are the most liquid.) In order for the index to reflect the most up-to-date corporate representation, and reflect accurately the credit risk environment, the composition of the index is re-examined frequently and firms are added or removed according to certain eligibility criteria. In particular, a new CDX series is introduced every six months on March 20 and September 20 of each year (or the next business day if the pre-specified day is on a weekend, or a holiday as designated by the New York Stock Exchange).

We perform a calibration exercise on the standardized tranche prices of the Series 7 CDX index (ticker symbol CDX.NA.IG.7) for October 31, 2006 as shown in the first two rows of Table 3.3. Also, for comparison purposes, we provide the fitting results on the Series 4 CDX index for August 23, 2004 (but without calibration to single name market data which were not able to obtain) and compare it against the results of Mortensen (2006).

Let us note that the credit quality of the CDX index has been changing with every new series that is introduced, as shown from the spread data in the first two rows of the Tables 3.3 and 3.5. As a result, both the overall level of the tranche spreads and their bid/ask spreads are decreasing, and for this reason we choose to illustrate the calibration results for both credit environments that are roughly two years apart.

3.3. Grouping Procedure. We apply the homogeneous group structure framework to assign the 125 components of the CDX index into separate groups that will describe the credit quality of the portfolio in more detail. We choose seven groups of firms with their sizes listed on the titles of the plots of Figure 3.1. The first group is comprised of the 8 firms with the smallest 5-year CDS spreads on October 31, 2006, the second group has the 10 firms with the next smallest spreads, and so on.

Figure 3.1 gives the breakdown of the 125 firms into the seven groups according to their 5-year CDS spread. The last two groups (groups 6 and 7) represent the above average credit-risk firms (when compared to the overall average of the CDX index), whereas the remaining firms are classified into five groups with relatively small default-risk. The default

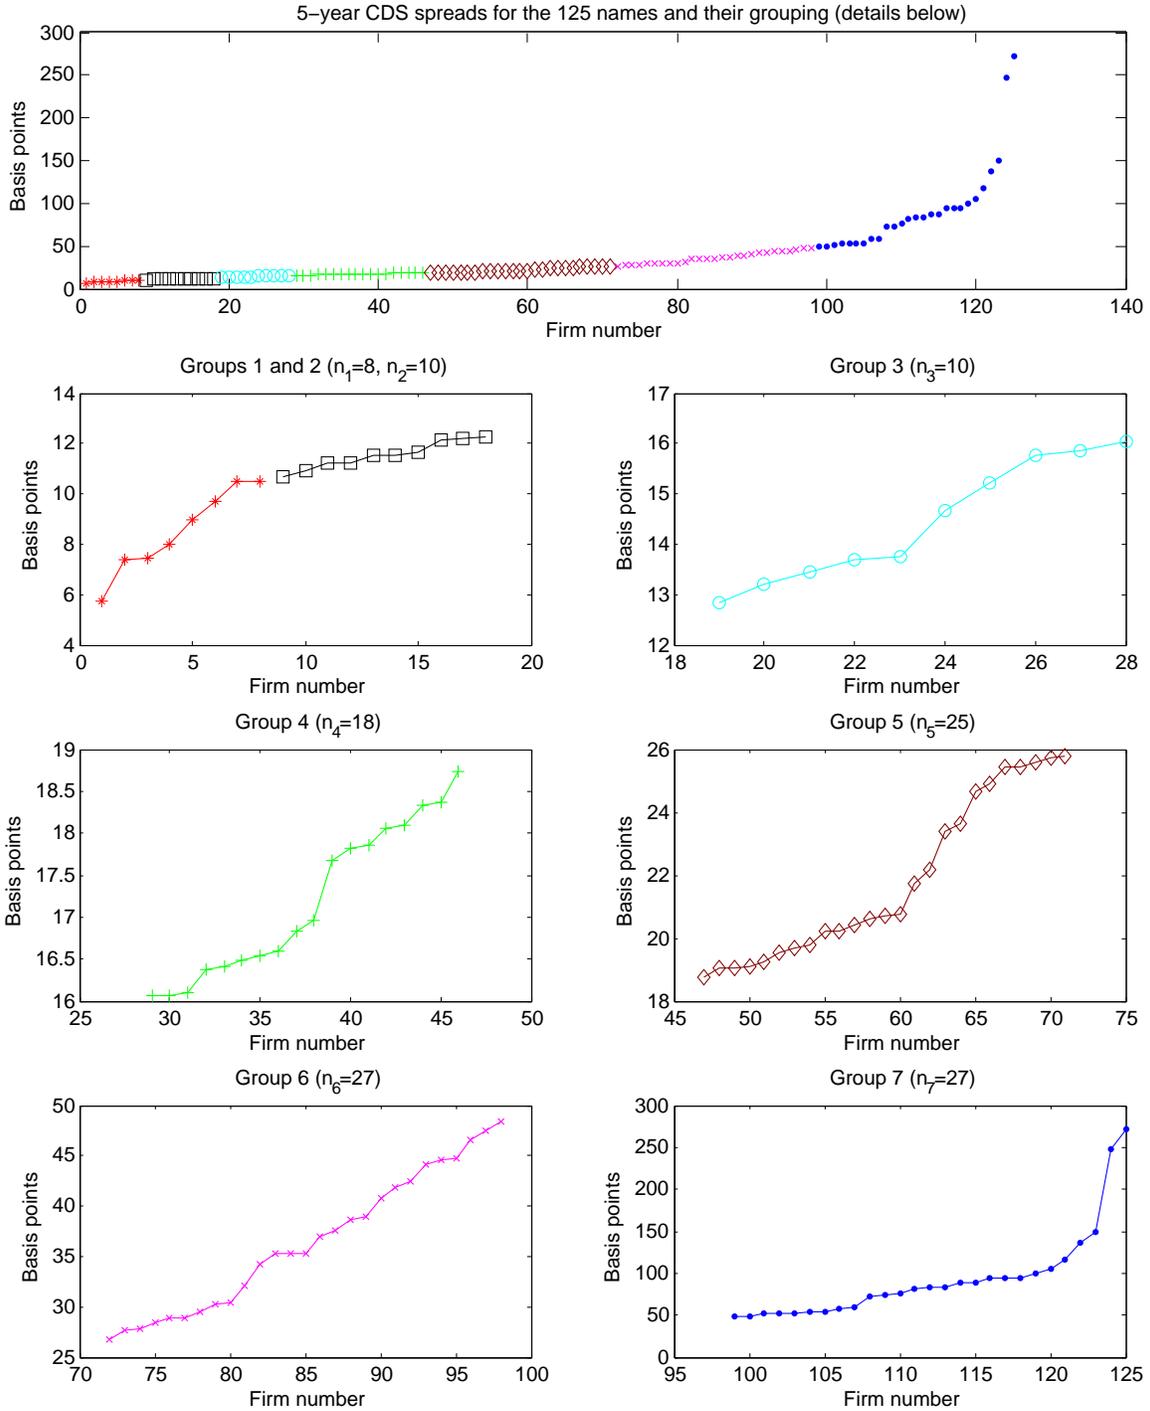


FIGURE 3.1. The 5-year CDS spreads of the 125 names of the CDX index on October 31, 2006 and their grouping.

intensity processes $\lambda^1, \dots, \lambda^7$ as given in (2.1) for the seven homogeneous groups will have term structure given by the average 1-, 3-, 5-, 7-, and 10-year CDS spreads of the firms within each group. These term structures on October 31, 2006 are given in Figure 3.2.

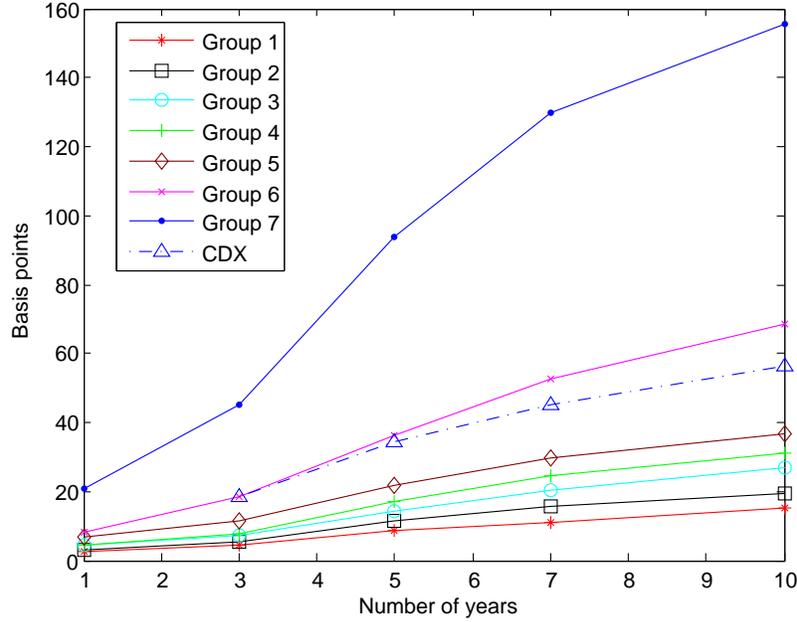


FIGURE 3.2. *CDS term structure of the CDX index (dashed line) and the average CDS term structure of the firms of the seven homogeneous groups on October 31, 2006. Source: Bloomberg Generic Database (CBGN).*

Another possible categorization of the 125 firms in groups is by industry sector but this does not reflect accurately the homogeneity assumption we make. Indeed, the firm with the widest 5-year spread on October 31 (The Gap Inc., 271 basis points) is in the same sector as one of the smallest-spread names (Jones Apparel Group, Inc., 10.5 basis points), which obviously do not seem to share the same default intensity process. Under the 5-year CDS level sorting criteria these two names are placed in the first and last groups.

TABLE 3.2. *The implied hazard curves for the seven groups and the CDX Series 7 index on October 31, 2006, assuming piecewise constant intensities.*

	Size	Year				
		1	3	5	7	10
Group 1	8	0.0017	0.0031	0.0065	0.0089	0.0139
Group 2	10	0.0020	0.0039	0.0088	0.0129	0.0177
Group 3	10	0.0028	0.0050	0.0111	0.0172	0.0252
Group 4	18	0.0028	0.0055	0.0133	0.0210	0.0297
Group 5	25	0.0045	0.0080	0.0169	0.0257	0.0350
Group 6	27	0.0054	0.0129	0.0291	0.0488	0.0746
Group 7	27	0.0135	0.0325	0.0846	0.1617	0.2561
CDX	125	-	0.0127	0.0272	0.0400	0.0570

Since the default intensity process λ^i is decomposed into an idiosyncratic, X^i , and a systematic risk factor, Z , we will calibrate the common market-risk factor against the CDX

index spread, i.e., the coupon that is paid for protection against the entire index (or the 0–100% tranche). The term structure is plotted in Figure 3.2 (dashed line). Notice that it is available only in 3-, 5-, 7-, and 10-year quotes.

We translate the average CDS term structure within each group into representative implied hazard rate curves for each group. For any firm in group i let $\bar{\tau}_i$ denote the representative default time (that is, a stopping time that shares the same law as the default time of any firm in the same group). The implied hazard rate of default at time t , $\lambda_f^i(t)$ (also known as the *forward intensity*), is defined by

$$\lambda_f^i(t) = -\frac{d}{dt} \log \mathbb{P}\{\bar{\tau}_i > t\}.$$

Assuming a piecewise constant hazard rate, we construct the implied hazard curve for each of the groups and the CDX index. The term structure is given in Table 3.2, and notice that groups 6 and 7 are considerably riskier than the rest consistent with the average CDS term structure in Figure 3.2. Furthermore, the overall CDX index on October 31, 2006 has a term structure between those of the fifth and sixth homogeneous groups.

3.4. Preliminary Fit to Data. There are 46 free parameters in the loss distribution expressions: $\alpha_1, \dots, \alpha_7, \bar{\sigma}_1^2, \dots, \bar{\sigma}_7^2, x_{0,1}, \dots, x_{0,7}, \bar{x}_1, \dots, \bar{x}_7, V_{1,1}, \dots, V_{1,7}, c_1, \dots, c_7, \alpha_z, \sigma_z^2, z_0, \bar{z}$, and only five tranche spreads, so our model is, obviously, over-parameterized. As ill-conditioned as this might look, it is rather a typical situation when a rich modeling framework is used for the pricing of basket credit derivatives. For instance, the total number of parameters to be estimated in the model proposed by Mortensen (2006) is $6 + 7n$ (before making the dimension reduction by considering equal parameters among the processes for different firms).

As such, this is not a disadvantage of the proposed setup because we explicitly allow for separate dynamics within each homogeneous group. Fundamentally, this is an improvement over the case where each underlying firm in the pool of the 125 names would have its own dynamics. By separating firms of *similar* credit risk (as measured by their 5-year CDS spread) and modeling their aggregate credit risk within their homogeneous group we obtain more accurate estimates of their default probabilities rather than assuming the existence of a unique specification of each firm’s credit risk.

In this section, we describe a rough calibration to the October 31, 2006 data. Later we describe a more ideal calibration procedure involving optimization, whose implementation is beyond the scope of the current paper. From the implied hazard curves of Table 3.2 we estimate the long-run levels $\bar{x}_1, \dots, \bar{x}_7$ of the seven intensity processes by dividing the 10-year hazard rate with the size of the group, i.e., $\hat{x}_i = \widehat{\lambda}_{10}^i/n_i$. This is because in the i th homogeneous group the “average” default intensity curve $t \mapsto \lambda_t^i$ represents the default intensity of the *first default time* τ^* among the n_i firm.

We estimate the long-run level of the common intensity factor, \bar{z} , from the 5-year hazard curve of the CDX index given in Table 3.2, and with this we estimate the “correlation” parameters c_1, \dots, c_7 of the idiosyncratic and market factors by $\hat{c}_i = \widehat{\hat{x}}_i/\widehat{\bar{z}}$. Such an estimator

tends to assign larger values to the higher default-risk groups because high-risk firms are more correlated to the overall state of the economy.

The starting values of the default intensity processes, $x_{0,1}, \dots, x_{0,7}, z_0$, are set equal to the estimated 1-year hazard rate. To obtain z_0 , since there is no data point for the 1-year quote of the CDX index we linearly extrapolate from the 3- and 5-year quotes of the hazard rate. Ideally, we would like to obtain the estimates of the starting values $x_{0,1}, \dots, x_{0,7}$ from the relationship $\widehat{\lambda}_1^i/n_i = \widehat{x}_{0,i} + \widehat{c}_i \widehat{z}_0$, but this is not feasible for the present dataset, as some of the resulting estimates are either too small, or even zero.

TABLE 3.3. *CDX tranche Series 7 (mid) spreads and bid/ask spreads on October 31, 2006, and Homogeneous Groups with Stochastic Volatility model (HGSV) fit. The equity tranche is quoted in percentage points of 3% (the detachment point) plus 500 bps running spread, while the other tranches are quoted in bps. The parameters are given in Table 3.4.*

	Tranche					RMSE
	0–3%	3–7%	7–10%	10–15%	15–30%	
Mid-market spread	23.66%	88.26	28.75	7.25	3.43	
Bid/ask spread	0.31%	1.25	1.11	0.64	0.97	
HGSV	23.76%	87.35	31.40	7.17	0.08	1.91

TABLE 3.4. *Parameters for the fit of Table 3.3. Note also that $\delta = 0.35$, and $r = 0.05$.*

	Group							Common
	1	2	3	4	5	6	7	
n	8	10	10	18	25	27	27	-
α	0.06	0.1	0.11	0.12	0.13	0.15	0.2	0.05
$\sqrt{\sigma^2}$	0.06	0.07	0.07	0.07	0.07	0.09	0.23	0.01
$V_1/10^{-4}$	-0.7	-1.07	-1.2	-1.43	1.43	1.35	1.65	-
$\bar{x}/10^{-3}$	2.1	2.2	2.9	2.1	1.8	3.2	9.9	2.72
$x_0/10^{-3}$	1.6	1.5	2.6	2.3	1.9	1.8	5.6	1.27
c	0.65	0.69	0.96	0.65	0.54	1.07	3.64	-

We fit the tranches with the remaining parameters, after setting the recovery rate, δ , equal to 0.35, and the risk-free interest rate, r , equal to 0.05. The results of the fitting are shown in Table 3.3 and the parameters used are listed in Table 3.4. The root-mean-squared error (RMSE) listed in the last column is calculated as (following Mortensen, 2006)

$$\text{RMSE} = \left(\frac{1}{5} \sum_{k=1}^5 \left(\frac{s_k - s_k^*}{s_k^{*,a} - s_k^{*,b}} \right)^2 \right)^{1/2},$$

where s_k is the spread of the k th tranche (or up-front fee in the case of the equity tranche) from our model, $s_k^{*,a}$ is the market ask spread, $s_k^{*,b}$ is the market bid spread, and $s_k^* = (s_k^{*,a} + s_k^{*,b})/2$

is the market mid spread. Other error measures may be used as well, but the RMSE conveys the idea of an overall good fit.

3.5. Iterative Calibration Procedure. We stress that the preceding calibration exercise follows a preliminary and partially *ad hoc* procedure to demonstrate the capability of the model. A better procedure we have begun to develop iterates between the multi-name tranche data and the single-name information. To elaborate, recall that in our model a typical firm in group i with default time $\bar{\tau}_i$ has intensity $\lambda^i = X^i + c_i Z$. The survival probability is approximated by

$$(3.1) \quad \mathbb{P}\{\bar{\tau}_i > T\} \approx A_i(T)e^{-B_i(T)x_0^i}[1 + \tanh(D_1^i(T)x_0^i + D_2^i(T))]A_z(T)e^{-B_z(T)z_0},$$

where A_i , B_i , D_1^i , and D_2^i are given by (2.11), (2.12), (2.27), and (2.28) respectively with the parameters corresponding to group i and $m = 1$, and A_z and B_z are given by (2.11) and (2.12) with the parameters of Z and $m = c_i$.

Given the Z parameters $(\alpha_z, \bar{z}, \sigma_z^2, z_0)$ and the mixing weights c_i we would fit the survival curve for the individual groups to the formula (3.1) to obtain estimates of $(\alpha_i, \bar{x}_i, \bar{\sigma}_i^2, x_{0,i}, V_{1,i})$. Then the Z parameters are adjusted to fit the tranches and the whole procedure is iterated until it stabilizes. This requires two optimization steps: the fit of the X^i -parameters in the groups and the fit of Z -parameters to the tranches. We do not attempt this here.

3.6. Effect of The Homogeneous-Group Setup. From the fitting results of Table 3.3 we see that the proposed setup with seven homogeneous groups and the calibrated parameters shown in Table 3.4 fits the equity and the three mezzanine tranches very well, each within one bid/ask spread deviation or less. The senior tranche (15–30%) is underpriced; the model computes a premium of less than one tenth of a basis point for protection against losses in that range, while the market quote was 3.43 basis points. Considering the bid/ask spread of one basis point, this deviation is rather large.

Past results regarding CDO pricing with intensity-based models (see Duffie and Gârleanu, 2001; Mortensen, 2006) have indicated that pure-diffusion models for the default intensity process are not able to generate loss distributions with the market-implied characteristics. This characteristic is simply the longer tails required for the probability mass function of L_t resulting from high-correlation situations.

However, we find that this remark applies only marginally to the present framework. Indeed, if that was the case the other tranches would have been considerably mispriced as well, since this is exactly the behavior of a pure-diffusion model as Mortensen (2006) notes in the calibration example (early peaks and thin tails of the loss distribution). Furthermore, the presence of the stochastic volatility correction term in the Laplace transform, required for the computation of the conditional probability mass function, introduces additional control over the probability density of the mid-range number of losses, but does not affect the probability of a small number of defaults.

The most important feature of the proposed framework is the categorization of the portfolio of firms into groups with homogeneous characteristics. As aggregate as it seems, modeling the firms' default intensities as a single process as long as they belong in the same homogeneous

group is *more accurate*, in terms of analyzing the overall credit risk of the portfolio, compared to some other pricing methodologies that are used in practice.

In applications of factor models to the pricing of CDOs, the inevitable simplifying assumption (or, compromise) is to consider that all firms' idiosyncratic components of the default intensity processes are, statistically, the same. This is done by setting all, or almost all, the parameters equal to each other and thus using the common ones to fit the tranche spreads to the market observed. The reason for this is to reduce the number of free parameters available for calibration, and hopefully use an optimization program for the minimization of an error measure.

The credit quality of the 125 firm-constituents of the CDX index, as measured by their 5-year CDS spread, can vary considerably (see top panel of Figure 3.1 for an example on October 31, 2006) and they should not be considered homogeneous. Dealer quotes for the tranches of CDX or iTraxx are, almost exclusively, computed by models having the assumption of a single idiosyncratic component among all the 125 processes. This mollifies the effect very high-risk firms have on the portfolio loss distribution, thus attributing the overall default risk evenly among all the firms, or, worse, to the systematic factor component of the default intensity. In turn, this attributes a larger-than-real weight of the default probabilities to the common market factor, which overestimates the probability of many defaults, which consequently prices the senior tranche very highly.

Therefore, the overly-simplistic assumption of extreme homogeneity of the idiosyncratic components of the default intensity explains, partially, the large premia of senior and super-senior tranches (i.e., the 15–30% and 30–100% tranches for CDX and the 12–22% and 22–100% tranches for iTraxx) we see in data. This is in accordance with Duffie and Gârleanu (2001) where, even though the asymmetrical information of the underlying assets is alleviated by the tranching of losses, there is still a considerable *liquidity premium* for the virtually improbable losses. This premium associated with the senior and super senior tranches is exogenous to the default premium dictated by pricing models. Sircar and Zariphopoulou (2006) reach the same conclusion via an alternative pricing approach using utility valuation.

We stress the fact that the breakdown of the 125 names in seven groups of variant credit risk explains the tranche spreads very closely, and is offered for the pricing of CDOs, and particularly bespoke CDOs, where a more inhomogeneous pool of firms (including speculative grade firms) is usually used.

We perform, also, a fitting example on the CDX tranches from August 23, 2004, and compare against the fitting of Mortensen (2006). The results are given in Table 3.5 and the parameters are listed in Table 3.6. Let us note that Mortensen (2006) uses the U.S. swap curve for the default-free interest rate, instead of our constant $r = 0.05$, but keeps $\delta = 0.35$.

It is not surprising that the model of homogeneous groups with stochastic volatility fits the equity and mezzanine tranches considerably better than all the other models, since it has many more free parameters that govern the dynamics. However, even after an exhaustive survey of the model's free parameters, the senior tranche remains underpriced which attests to the aforementioned argument of the existence of an additional premium included in the senior tranche spread. On the other hand, the jump-diffusion intensities model does fit the

TABLE 3.5. *Market spreads for the CDX Series 4 index on August 23, 2004, and fit comparisons between the Homogeneous Groups with Stochastic Volatility model (HGSV) and five other popular models. The equity tranche is quoted in percentage points and all the other tranches in bps. The parameters are given in Table 3.6.*

	Tranche					RMSE
	0–3%	3–7%	7–10%	10–15%	15–30%	
Mid-market spread	40%	312.5	122.5	42.5	12.5	
Bid/ask spread	2%	15	7	7	3	
HGSV	40.8%	311.5	123.4	40.4	1.6	1.64
Jump-diffusion intensities	46.9%	340.2	119.7	61.9	14.3	2.17
Pure-diffusion intensities	49.3%	442.9	94.9	16.8	0.4	5.34
Gaussian copula	46.8%	474.4	131.8	36.9	2.9	5.3
RFL Gaussian copula	48.6%	334.9	125.5	66.5	9.2	2.59
Double- t copula	45.1%	367.0	114.9	54.9	20	2.44

Source: Last five rows Mortensen (2006).

TABLE 3.6. *Parameters for the fit of Table 3.5. Note also that $\delta = 0.35$, and $r = 0.05$.*

	Group							Common
	1	2	3	4	5	6	7	
n	8	10	10	19	25	26	27	-
α	0.08	0.09	0.1	0.21	0.23	0.165	0.17	0.05
$\sqrt{\sigma^2}$	0.06	0.07	0.07	0.18	0.2	0.41	0.45	0.01
$V_1/10^{-4}$	-0.07	-1.07	-1.2	14.5	14.43	14.5	20.0	-
$\bar{x}/10^{-3}$	4.5	4.7	5.9	6.2	8.0	10.5	10.6	2.72
$x_0/10^{-3}$	3.42	5.5	5.7	5.5	6.5	8.0	14.9	1.27
c	0.65	0.69	0.96	0.65	0.54	1.07	3.49	-

senior tranche quite well, but at the cost of grossly over-estimating the third mezzanine tranche and, to a lesser extent, the first mezzanine tranche also.

3.7. Effect of the Grouping and Stochastic Volatility on the Tranches. The number of groups and the number of firms per group are directly associated with the thickness of the tail of the loss distribution. Due to the required computation of the conditional probability mass function (2.33) for every group, we limit the size of each group to less than or equal to 30, in order to avoid the large numerical error that is inherited from the Euler-Maclaurin alternating sum. Therefore, we require at least five homogeneous groups for a 125-firm portfolio such as the pools of the CDX and iTraxx indices. In practice, we need around seven groups, at least one of which needs to have large variance and large positive V_1 parameter in order to extend the tail of the distribution.

The effect of the fast mean-reverting stochastic volatility aggregate parameter is to move mass within the loss distribution and alter its tail. Notice that in the calibrated results

both positive and negative values of the aggregate parameter are required to control the middle part of the portfolio mass function efficiently. (See also Figure 2.1 for the effect of the aggregate parameter on the single-group loss distribution.)

3.8. 2008 Credit Crisis. Since the first drafts of the present paper, there has been an enormous crisis in the world economy, and in particular, the credit derivatives market. The origination of the problem for financial institutions is in their over-exposure to mortgage-backed securities, and CDOs written on those. However, the market in corporate credit, especially investment grade which we consider here empirically, has seen spreads widen massively. We present in this section, a preliminary fitting to CDX tranche data from October 16, 2008, after the failure of Lehman Brothers, and the bailouts of Fannie Mae and Freddie Mac, and the takeover of Washington Mutual. The last three names were part of the CDX IG Series 10 index at that time, and as a result the CDO tranches now reference a pool of only 122 firms instead of the usual 125.

Table 3.7 presents the results of a rough fit where we divided the 122 firms into seven groups as described in Table 3.8. During this period of tremendous risk reduction and very poor liquidity conditions the tranches of the CDX index were not as actively quoted. Our preliminary fitting shows the HGSV model is able to get within 10–21% of the quoted spreads of the equity and the three mezzanine of the CDX index by (i) increasing the contribution of the systemic risk factor Z by increasing the c_i parameters, and (ii) fine tuning with the stochastic volatility parameters V_1 . A more detailed analysis of this period of unprecedented credit conditions remains for future work. (It is interesting to note that in early November 2008 the first mezzanine tranche (3–7%) joined the equity tranche in being quoted with upfront points plus a running spread, instead of the usual spread convention, illustrating the extreme counterparty risk aversion among primary dealers.)

TABLE 3.7. *CDX tranche Series 10 (mid) spreads on October 16, 2008, and Homogeneous Groups with Stochastic Volatility model (HGSV) rough preliminary fit. The equity tranche is quoted in percentage points of 3% (the detachment point) plus 500 bps running spread, while the other tranches are quoted in bps.*

	Tranche				
	0–3%	3–7%	7–10%	10–15%	15–30%
Mid-market spread	78.25%	1686.75	879.10	343.81	79.53
HGSV	88.03%	2014.23	686.58	290.48	32.29

TABLE 3.8. *Group structure of the October 16, 2008 CDX Series 10.*

	Group						
	1	2	3	4	5	6	7
Size	5	10	12	21	21	27	26
Avg. 5-Yr CDS spread	37.6	55.8	72.0	93.4	139.8	246.0	813.2

4. CONCLUSIONS

We present a new method for the computation of the loss distribution for a large pool of firms which is necessary for the pricing of CDOs and variants of this type. In particular, we introduce two practical implementations for the derivation of the loss distribution in a common factor model framework. Firstly, we allow for the categorization of the firms into groups of identical default intensity processes. This describes the granularity of the overall credit risk of the portfolio in more detail as we argued in Section 3.6. Secondly, we offer an asymptotic approximation for the loss distribution by incorporating a fast mean-reverting stochastic volatility process in the default intensity process. We do so in order to produce a similar effect on the loss distribution as do the jump-diffusion models of Duffie and Gârleanu (2001) and Mortensen (2006)—namely shift the mass of the loss distribution towards the right. However, unlike the aforementioned sources, the asymptotic approximation terms are given in closed form, thus reducing considerably the computational burden associated with the final Fourier transforms.

One practical issue plaguing the calibration of CDO models, is that rarely are complete data regarding all 125 constituent firms of the CDX index available. The homogeneous group assumption mitigates the incomplete CDS term structure that a considerable amount of the firms have, and instead allows for an “average” and representative CDS term structure to be allocated to each of the groups.

Let us note also that once a model has been calibrated against a set of data, such as the CDX index, the fitted parameters may be used for the pricing of other derivatives that use the same parameter information, such as CDOs or bespoke CDOs. In particular, if the firm pool of a CDO has common constituents as the CDX index firm pool, then with appropriate groupings we can use the fitted parameters from the processes of the CDX index as estimates for the parameters for the default intensity processes of their CDO counterparts.

Additionally, most of the proposed pricing models in existence make simplifying assumptions regarding the idiosyncratic components of the default intensity process of each firm, usually by making the idiosyncratic parts of these processes have the same parameters, in order to reduce the number of parameters that need to be estimated from data. As shown in Mortensen (2006), this leads to the overestimation of the importance of the systematic component in the default intensity, and subsequently to the mispricing of tranches, and especially the senior one. The homogeneous group setup illustrates the consistent overpricing of the senior tranches, and indicates the existence of a large liquidity premium.

The model presented can be seen also as being between the top-down and the bottom-up credit pricing frameworks, since the single-name components are treated only as aggregate inputs in the specification of the loss distribution. The focus of present problems in multiname credit derivatives seems to be shifting towards the pricing of options on CDOs and other exotic derivatives. This involves more elaborate characterization of the joint distribution of default times, and not just the convolution of the marginal conditional loss distribution given the systematic factor that we used above. The methodology to address the pricing and hedging of these complex derivatives is still largely in development—mainly using top-down

models—and it remains to be seen how well the present bottom-up models will integrate with it.

APPENDIX A. FOURIER TRANSFORM DETAILS

We provide below a quick review of the Fourier inversions we performed for the model calibration. We refer to Papageorgiou (2007) for an analytical description of the application of the fast Fourier transform on the inversion of the characteristic and generating functions.

A.1. Inversion of The Characteristic Function. Fix a time $t > 0$. If we define f_U to be the density function of the random variable $U_t = \int_0^t Z_s ds$, i.e.,

$$\mathbb{P}\{U_t \in dx\} = f_U(x) dx, \quad x > 0,$$

then its *characteristic function* φ (or *Fourier transform*) is given by

$$\varphi(\eta) := \mathbb{E}[e^{i\eta U_t}] = \int_0^\infty e^{i\eta x} \mathbb{P}\{U_t \in dx\}$$

where η is a real number and i is the imaginary unit. Since the functions A and B in (2.11) and (2.12) of the Laplace transform (2.8) are analytic, the characteristic function is given by

$$\varphi(\eta) = \left(\frac{2\gamma_z e^{(\alpha_z + \gamma_z)t/2}}{(\alpha_z + \gamma_z)(e^{\gamma_z t} - 1) + 2\gamma_z} \right)^{2\alpha_z \bar{z} / \sigma_z^2} \exp \left\{ - \frac{2i\eta(e^{\gamma_z t} - 1)z}{(\alpha_z + \gamma_z)(e^{\gamma_z t} - 1) + 2\gamma_z} \right\}$$

where $\gamma_z = \sqrt{\alpha_z^2 - 2i\eta\sigma_z^2}$.

Since f_U is real and positive, φ is odd in its imaginary part and even in its real, thus f_U is written in terms of φ as

$$f_U(x) = \frac{1}{2\pi} \int_{\mathbb{R}} e^{-i\eta x} \varphi(\eta) d\eta = \frac{1}{\pi} \int_0^\infty \operatorname{Re}(e^{-i\eta x} \varphi(\eta)) d\eta.$$

Using the equivalent discrete Fourier transform we evaluate numerically the latter integral on the grid $\{x_k; k = 1, \dots, M\}$, where $x_k = (k-1)\Delta x$ for some positive number Δx and large number M . The values of f_U on this grid are

$$f_U(x_k) = \frac{1}{\pi} \sum_{j=1}^M \operatorname{Re} \left(e^{-i\Delta\eta \Delta x (j-1)(k-1)} \varphi(\Delta\eta(j-1)\Delta\eta) \right), \quad k = 1, \dots, M, \quad \Delta x \cdot \Delta\eta = \frac{2\pi}{M}.$$

We evaluate this sum by fast Fourier transform methods similar to those implemented by Carr and Madan (1999).

Remark A.1. Let us note that since the random variable U_t is positive, we could consider a Laplace transform inversion instead of the Fourier transform, since it is well defined. We choose the latter though due to the wider availability of inversion algorithms.

A.2. Inversion of The Generating Function. Fix $t > 0$ and let us define $g : (0, 1) \times \mathbb{R}_+ \rightarrow (0, 1)$ to be the conditional generating function of L_t given the common integrated factor U_t with

$$g(z, v) := \mathbb{E} [z^{L_t} | U_t = v],$$

which is computed by (2.4). Then, given g the conditional loss distribution is computed by the contour integral

$$\mathbb{P}\{L_t = k | U_t = v\} = \frac{1}{2\pi i} \oint_C \frac{g(z, v)}{z^{k+1}} dz, \quad k = 0, \dots, M-1,$$

where i is the imaginary unit, and C is a closed curve that encloses all the poles of $g(\cdot, v)$ for $v \in \mathbb{R}_+$. Following Cavers (1978), we set C to be a circle of unit radius (because $z \in (0, 1)$) which simplifies the contour integral to

$$\mathbb{P}\{L_t = k | U_t = v\} = \frac{1}{2\pi} \int_0^{2\pi} \operatorname{Re} \left(e^{ikz} g(e^{iz}, v) \right) dz, \quad k = 0, \dots, M-1.$$

Adopting a discretization scheme similar to that for the characteristic function inversion in Section A.1 above we can write

$$\mathbb{P}\{L_t = k | U_t = v\} = \frac{1}{M} \sum_{j=1}^M \operatorname{Re} \left(e^{-i \frac{2\pi}{M}(j-1)k} g \left(e^{i \frac{2\pi}{M}(j-1)}, v \right) \right), \quad k = 0, \dots, M-1,$$

and a fast Fourier transformation yields the numerical results. Of course, M has to be greater than or equal to $n+1$ in order for the values of the FFT to correspond to the correct conditional probabilities.

REFERENCES

- Andersen, L. and Sidenius, J. (2005). CDO pricing with factor models: survey and comments. *Journal of Credit Risk*, 1(3):71–88.
- Andersen, L., Sidenius, J., and Basu, S. (2003). All your hedges in one basket. *Risk*, November:67–72.
- Arnsdorf, M. and Halperin, I. (2007). BSLP: Markovian bivariate spread-loss model for portfolio credit derivatives. Working paper, JP Morgan.
- Bielecki, T. R., Vidozzi, A., and Vidozzi, L. (2006). An efficient approach to valuation of credit basket products and ratings triggered step-up bonds. Working paper, Illinois Institute of Technology.
- Brigo, D. and Mercurio, F. (2006). *Interest Rate Models: Theory and Practice*. Springer, 2nd edition.
- Carr, P. and Madan, D. B. (1999). Option valuation using the fast fourier transform. *Journal of Computational Finance*, 2(4):61–73.
- Cavers, J. K. (1978). On the fast Fourier transform inversion of probability generating functions. *IMA Journal of Applied Mathematics*, 22(3):275–282.
- Cherubini, U., Luciano, E., and Vecchiato, W. (2004). *Copula Methods in Finance*. Wiley.
- Cotton, P., Fouque, J.-P., Papanicolaou, G., and Sircar, R. (2004). Stochastic volatility corrections for interest rate derivatives. *Mathematical Finance*, 14(2):173–200.

- Credit Suisse First Boston (1997). CreditRisk⁺: A credit risk management framework. Technical report.
- Di Graziano, G. and Rogers, L. (2006). A dynamic approach to the modelling of correlation credit derivatives using Markov chains. Working paper, University of Cambridge.
- Duffie, D. and Gârleanu, N. (2001). Risk and valuation of collateralized debt obligations. *Financial Analysts Journal*, 57(1):41–59.
- Errais, E., Giesecke, K., and Goldberg, L. R. (2007). Pricing credit from the top down with affine point processes. Working paper, Stanford University.
- Feller, W. (1951). Two singular diffusion problems. *The Annals of Mathematics*, 54(1):173–182.
- Fouque, J.-P., Papanicolaou, G., and Sircar, R. (2000). *Derivatives in Financial Markets with Stochastic Volatility*. Cambridge University Press.
- Fouque, J.-P., Papanicolaou, G., Sircar, R., and Sølna, K. (2003). Singular perturbations in option pricing. *SIAM Journal of Applied Mathematics*, 63(5):1648–1665.
- Fouque, J.-P., Wignall, B. C., and Zhou, X. (2008). Modeling correlated defaults: First passage model under stochastic volatility. *Journal of Computational Finance*. to appear.
- Fouque, J.-P. and Zhou, X. (2006). Perturbed Gaussian copula. To appear in *Advances in Econometrics*.
- Frey, R. and Backhaus, J. (2006). Credit derivatives in models with interacting default intensities: a Markovian approach. Working paper, University of Leipzig.
- Frey, R. and McNeil, A. J. (2003). Dependent defaults in models of portfolio credit risk. *Journal of Risk*, 6(1):59–92.
- Gupton, G. M., Finger, C. C., and Bhatia, M. (1997). CreditMetrics. Technical report, RiskMetrics Group.
- Hull, J. and White, A. (2004). Valuation of a CDO and an n-th to default CDS without Monte Carlo simulation. *Journal of Derivatives*, 12(2):8–23.
- Hurd, T. R. and Kuznetsov, A. (2005). Fast CDO computations in the affine Markov chain model. Working paper, McMaster University.
- Lando, D. (2004). *Credit Risk Modeling: Theory and Applications*. Princeton University Press.
- Li, D. X. (2000). On default correlation: A copula function approach. *Journal of Fixed Income*, 9(4):43–54.
- Mortensen, A. (2006). Semi-analytical valuation of basket credit derivatives in intensity-based models. *Journal of Derivatives*, 13(4):8–26.
- Musiela, M. and Rutkowski, M. (2007). *Martingale Methods in Financial Modelling*. Springer, 2nd edition.
- Papageorgiou, E. (2007). *Single-name and Multi-name Credit Derivatives: Pricing and Calibration Using Multiscale Asymptotic Methods*. PhD thesis, Princeton University.
- Papageorgiou, E. and Sircar, R. (2008). Multiscale intensity models for single name credit derivatives. *Applied Mathematical Finance*, 15(1):73–105.
- Schönbucher, P. J. (2003). *Credit Derivatives Pricing Models: Models, Pricing and Implementation*. Wiley.

Sircar, R. and Zariphopoulou, T. (2006). Utility valuation of credit derivatives and application to CDOs. Working paper, Princeton University.

(Both authors) DEPARTMENT OF OPERATIONS RESEARCH & FINANCIAL ENGINEERING, PRINCETON UNIVERSITY, E-QUAD, PRINCETON, NJ 08544

E-mail address, Evan Papageorgiou: evanp@princeton.edu

URL: <http://www.princeton.edu/~evanp>

E-mail address, Ronnie Sircar: sircar@princeton.edu

URL: <http://www.princeton.edu/~sircar>