Two-Dimensional Markovian Model for Dynamics of Aggregate Credit Loss

A.V. Lopatin and T. Misirpashaev
NumeriX Quantitative Research*

First version: February 16, 2007

This version: August 29, 2007

Abstract

We propose a new model for the dynamics of the aggregate credit portfolio loss. The model is Markovian in two dimensions with the state variables being the total accumulated loss L_t and the stochastic default intensity λ_t. The dynamics of the default intensity are governed by the equation $d\lambda_t = \kappa(\rho(L_t, t) - \lambda_t)dt + \sigma\sqrt{\lambda_t}dW_t$. The function ρ depends both on time t and accumulated loss L_t, providing sufficient freedom to calibrate the model to a generic distribution of loss. We develop a computationally efficient method for model calibration to the market of synthetic single tranche CDOs. The method is based on the Markovian projection technique which reduces the full model to a one-step Markov chain having the same marginal distributions of loss. We show that once the intensity function of the effective Markov chain consistent with the loss distribution implied by the tranches is found, the function ρ can be recovered with a very moderate computational effort. Because our model is Markovian and has low dimensionality, it offers a convenient framework for the pricing of dynamic credit instruments, such as options on indices and tranches, by backward induction. We calibrate the model to a set of recent market quotes on CDX index tranches and apply it to the pricing of tranche options.

*NumeriX LLC, 757 Third Avenue, New York, NY 10017

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

Two-Dimensional Markovian Model for Dynamics of Aggregate Credit Loss
Copyright ©1996–2007 by NumeriX LLC. All rights reserved.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system or transmitted in any form or by any means (electronic, mechanical, photocopying, recording or otherwise), or for any purpose, without the express written permission of NumeriX LLC.

Requests for permission or further information should be addressed to NumeriX LLC, 757 Third Avenue, 20th Floor, New York, NY 10017.

Numerix® is a registered trademark of NumeriX LLC. Brand and product names not associated with NumeriX LLC are trademarks or registered trademarks of their respective companies.

Corporate Headquarters
NumeriX LLC
757 Third Avenue, 20th Floor
New York, NY 10017
Tel: +1 212 302 2220
Fax: +1 212 302 6934

EMEA
NumeriX Software Ltd.
41 Eastcheap, 2nd Floor
London EC3M 1DT
Tel: + 44 20 7648 6100
Fax: + 44 20 7648 6139

Asia
NumeriX Japan Co., Ltd.
Tokyo Ginko Kyokai Building, 15th Floor
The Executive Center
1-3-1 Marunouchi, Chiyoda-ku
Tokyo 100-0005
Tel: 81-3-3216-7250
Fax: 81-3-3216-7289

Support
US & Canada
Tel: +1 877 Numeri0(686-3740)

Global
Tel: +1 212 302 0465
Fax: +1 212 302 6934

Customers with questions about our products or their license keys should always feel free to contact NumeriX technical support. E-mail: **support@numerix.com**

1 Introduction

Synthetic collateralized debt obligations (CDOs) are derivatives of the aggregate loss sustained by the seller of the protection on a portfolio of credit default swaps. The majority of standard CDO tranches can be statically replicated by a set of long and short positions in more elementary instruments called stop-loss options.[1] The payoff from a stop-loss option with maturity t and strike X is equal to $\max(L_t - X, 0)$, where L_t is the loss accumulated in the underlying portfolio by the maturity time t.

This replication is not directly useful for hedging purposes because standalone stop-loss options are not currently traded. However, it is extremely useful in model-based valuation because of the simplicity of stop-loss options which only depend on the distribution of loss at a single time horizon. An important consequence is that the value of an instrument replicated by a portfolio of stop-loss options only depends on the one-dimensional marginal distributions of the loss process and is insensitive to the dynamics of the temporal loss evolution. Therefore, it is possible to construct viable valuation models for synthetic CDO tranches by focusing solely on producing correct distributions of loss on a grid of relevant time horizons and ignoring the implied dynamics (or even leaving the dynamics undefined). Such models are often referred to as static. Most static models in active use today belong to the framework of factor models, reviewed by Andersen and Sidenius (2005b).

There are two main practical reasons to go beyond the static models. First, there are instruments that do not admit a replication by stop-loss options. These include forward-starting CDO tranches, options on tranches, leveraged super-senior tranches, and other innovative structured products such as constant proportion portfolio insurance (CPPI) and constant proportion debt obligations (CPDO). The second reason is the ambiguity of dynamic hedging and the difficulty of managing the risk of forward exposures on the basis of static models, even for positions in standard tranches.

While the potential of the new generation of factor models to build adequate dynamics of portfolio loss starting from loss distributions of individual obligors is certainly far from exhausted (see, e.g., Andersen, 2006, and Chapovsky et al., 2006), there is growing appreciation of the benefits of the direct modeling of the loss process L_t. The general framework of aggregate-loss-based approaches to basket credit derivatives was put forward by Giesecke and Goldberg (2005), Schönbucher (2005), and Sidenius et al. (2005). Examples of specific models can be found in the works by Bennani (2005), Brigo et al. (2006), Errais et al. (2006), Ding et al. (2006), and Longstaff and Rajan (2006).

Both Schönbucher (2005) and Sidenius et al. (2005) aimed to build a credit portfolio counterpart of the Heath-Jarrow-Morton (HJM) framework of the interest rate modeling. In the HJM framework, the problem of fitting the initial term structure of interest rates is non-existent because the discount curve serves as an initial condition and not as a calibration constraint. In the calibration of credit portfolio models, the role of the discount curve is played by the surface of the loss distribution, $\pi(L,t) = \mathrm{P}[L_t \leq L]$. In the spirit of the HJM framework, Schönbucher (2005) and Sidenius et al. (2005) eliminated the problem of the calibration to the distribution of loss by making it an initial condition. This, however, was achieved at a price of losing the ability to simulate the loss process without introducing a large number of additional stochastic degrees of freedom, which led to severe computational problems and accentuated the need for a more specific approach.

[1] The replication of super-senior tranches requires a similar set of options on recovery also; see Appendix A.

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

While the HJM framework indeed provides ultimate flexibility in fitting the market, many of the short rate models developed before HJM are also capable of fitting the entire discount curve. The flexibility of the calibration was achieved due to the presence of a free function of time in the drift term of the model-defining stochastic differential equation (SDE). The models developed within this scheme had a tremendous impact on the field and are still highly popular among practitioners. In view of this success, it appears reasonable to try an adaptation of the framework of short rates to the problem of credit portfolio loss.

As was pointed out by Schönbucher (2005), models based on an explicit, short-rate-like modeling of the loss intensity generally run into a problem with the calibration to the distribution of loss. Indeed, fitting an entire two-dimensional surface of loss can require a large number of free calibration parameters and is likely to be computationally burdensome. It might be argued that the information about the loss distribution coming from the standard tranches is too sparse to define a complete surface of loss. Brigo et al. (2006), Errais et al. (2006), Ding et al. (2006), and Longstaff and Rajan (2006) reported successful calibration to the tranche market. Their models are formulated in an open-ended way, so that it might in principle be possible to calibrate increasingly many tranches by adding new terms to the defining equations. However, the problem of finding a model based on an explicit equation for the loss intensity, and amenable to a reasonably fast calibration to a generic distribution of loss, has remained unresolved.[2]

In this work, we introduce a new two-dimensional intensity-based Markovian model of the aggregate credit loss. This model can be easily calibrated to a generic distribution of portfolio loss without sacrificing tractability and robustness. The calibration procedure consists of two steps. On the first step, we find the intensity of an auxiliary one-step Markov chain model consistent with the CDO tranches. Because the intensity of the Markov chain is a deterministic function of accumulated loss and time, it can be called *local* intensity, to distinguish it from the stochastic intensity of the full model. On the second step, the full two-dimensional model is calibrated to match the local intensity. The idea of exploring the link between the local intensity and the stochastic intensity is borrowed from the method of Markovian projection used by Dupire (1994) and Piterbarg (2006) to relate the local volatility and the stochastic volatility. For the purpose of credit applications, we extended the original formulation of the Markovian projection given by Gyöngy (1986) from diffusions to jump processes.

A model calibrated to the market quotes on CDO tranches can be used to price more complicated dynamic instruments. In this paper, we consider an application to the tranche option which, as we show, can be evaluated easily using the backward induction technique. In a numerical example, we calibrated our model to a set of recent quotes for the tranches on Dow Jones CDX.NA.IG.7 and calculated the values of the tranche option at different strikes.

The rest of the paper is organized as follows. In Section 2, we define our model and give a general discussion of its properties.

Section 3 is devoted to model calibration. In Section 3.1, we assume (unrealistically) that a full surface of loss distribution is available. This would be the case if we had arbitrage-free quotes for stop-loss options at all strikes and maturities. We show how to build an efficient numerical procedure for the calibration of the two-dimensional Markovian model once we know the local intensity of the auxiliary one-step Markov chain model. We show, furthermore, that finding the local intensity from a complete loss distribution is straightforward. In practice, only

[2] A solution alternative to ours was independently developed by Arnsdorf and Halperin (2007) and released when the first revision of our paper was in circulation.

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

a handful of quotes for CDO tranches with a sparse set of attachment points and maturities are available, so that it is not possible to restore the full distribution of loss without interpolation assumptions. We address this issue in Section 3.2. Instead of interpolating the loss, we choose to do a parametric fit for the coefficients in a specific analytical form for the local intensity. Numerical results for the calibration are given in Section 3.3.

Applications to dynamic pricing problems are discussed in Section 4. We begin by describing the general backward induction set-up in Section 4.1, which is followed by a discussion of numerical results for tranche options in Section 4.2.

We conclude in Section 5.

The Appendix consists of three parts completing the main text. Part A describes the cash-flows of the single tranche CDO and explains its replication by a portfolio of stop-loss options. Part B contains a digression into the method of Markovian projection for stochastic volatility modeling and our extension of Gyöngy's lemma for jump processes. Part C gives the details of the discretization scheme used in the numerical implementation.

2 The model

We work in the top-down framework and model the loss as an intensity-based process (see, e.g., Duffie, 2005, for general properties and definitions, and Giesecke and Goldberg, 2005, for a conceptual description of the top-down framework). Other possibilities for introducing a compact formulaic description of the loss process were tried. For example, Bennani (2005) postulates an SDE on the loss process without introducing an intensity. However, this sacrifices the discrete nature of default arrivals, which is generally considered an important feature to keep.

The minimal number of independent variables necessary to describe a state of an intensity-based process is two, the accumulated loss L_t and the intensity λ_t. We do not introduce any additional variables and postulate the following dynamics for the intensity,

$$d\lambda_t = \kappa\left(\rho(L_t, t) - \lambda_t\right) dt + \sigma\sqrt{\lambda_t} dW_t, \tag{1}$$

where W_t is the standard Brownian motion. We follow the work of Errais et al. (2006) in allowing for a back action of the loss onto the intensity process (thus going beyond the class of Cox processes which exclude such an action). We found, however, that restricting the model to the affine framework limits its ability to achieve a stable calibration to the market. The function $\rho(L_t, t)$, in general, is not linear in the accumulated loss L_t, and therefore our model is not affine. This function serves to provide sufficient freedom to calibrate to a generic distribution of loss $\pi(L, t)$. In contrast to the affine set-up of Errais et al. (2006), our model has no transform-based analytical solution for the stop-loss option. We will show, nevertheless, that an efficient numerical calibration to an arbitrary distribution of loss is possible.

Throughout this paper, we assume that the value of loss-given-default (LGD) is equal to the same amount h for all assets in the basket. Many authors, including Brigo et al. (2006) and Errais et al. (2006), point out the importance of a non-trivial LGD distribution for the market-matching ability of their models. We believe that our model can describe the market data even in its simplest form, with a deterministic LGD, because sufficient flexibility is already built-in via the function $\rho(L_t, t)$. Our model can be generalized to include stochastic LGD at

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

the cost of introducing the third dimension, which in our numerical experiments reduced the performance without a significant functional improvement.

Note that the calibration to the loss distribution will be achieved only by adjusting the function $\rho(L_t, t)$ in the drift term, but not the multiplier κ or the volatility σ in the diffusion term of Eq. (1). The volatility term is kept available to tune the dynamics of the model. Given the potential for the growth in the variety and liquidity of dynamics-sensitive credit instruments, we can envisage a scenario where the volatility will be calibrated to simpler instruments (for example, European tranche options) and then used to value more complex ones (for example, Bermudan tranche options). If necessary, a constant volatility can be generalized to a term structure. This is similar to the calibration strategy for the classic short rate models of interest rates, including Hull-White (HW), Black-Karasinski (BK), and Cox-Ingersoll-Ross (CIR). For these models, the term structure of volatilities is fitted to the options in a cycle of solver iterations, with the free function of time in the drift term being fitted to the discount curve inside each iteration.

We chose CIR-like dynamics (1) for the intensity only as a starting point. Similar models are possible based on single-factor or multi-factor BK-like equations. It is also possible to introduce jump terms in the intensity. The procedures described in this paper will remain applicable provided the model has a free function of loss and time in the drift term and does not lead to negative values of intensity. In the case of CIR-like dynamics (1) the flow of intensities through the boundary $\lambda = 0$ to the negative values is avoided as long as $\rho(L, t) > 0$.

3 Calibration

For a given function $\rho(L, t)$, the model defined by Eq. (1) can be easily treated numerically. Depending on the financial instrument, it can be solved either by a numerical integration of the backward Kolmogorov equation (as discussed in detail in Section 4.1) or by a forward Monte Carlo simulation. However, a direct iterative calibration would certainly be too time-consuming. In this Section we develop a computationally efficient procedure for the calibration to CDO tranches, which avoids massive iterative fitting.

The goal of the calibration is to find the function $\rho(L, t)$ consistent with the information about the loss distribution available from the market data. As mentioned in the previous section, the reversion strength κ and the volatility σ are not subject to calibration at this stage. We want to keep them unconstrained by the static information and potentially available for the calibration to dynamically sensitive instruments, such as options on tranches.

The starting point of our calibration procedure is the forward Kolmogorov equation for the joint density $p(\lambda, L, t)$ of intensity λ and loss L, following from Eq. (1),

$$\frac{\partial p(\lambda, L, t)}{\partial t} = \left(-\kappa \frac{\partial}{\partial \lambda} (\rho(L, t) - \lambda) + \frac{1}{2} \frac{\partial^2}{\partial \lambda^2} \sigma^2 \lambda \right) p(\lambda, L, t) + \lambda \left(\mathbf{1}_{L \geq h} p(\lambda, L - h, t) - p(\lambda, L, t) \right). \quad (2)$$

Here λ is restricted to be non-negative, and the loss L takes the values $L = 0, h, 2h, \ldots, N_{\max} h$, where $N_{\max}$ is the number of assets in the basket. The absence of the probability flow through the boundary $\lambda = 0$ to the negative intensities is ensured by the boundary condition

$$\left(-\kappa(\rho(L, t) - \lambda) + \frac{1}{2} \frac{\partial}{\partial \lambda} \sigma^2 \lambda \right) p(\lambda, L, t) = 0, \quad \lambda \to 0. \quad (3)$$

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

We also need an initial condition to Eq. (2), which obviously has the following form

$$p(\lambda, L, 0) = p_0(\lambda) \cdot \mathbf{1}_{L=0}. \tag{4}$$

The function $p_0(\lambda)$ is not fully fixed by the market because we cannot observe the distribution of default intensity. The choice of this function will be discussed in Section 3.2.

The probability density of loss is obtained by integrating the joint density over λ,

$$P(L,t) = \int_0^\infty p(\lambda, L, t) d\lambda. \tag{5}$$

3.1 Calibration to local intensity

In this section, we assume that the calibration target is the entire set of one-dimensional marginal loss densities, that is, the dependence $P(L,t) = \mathrm{P}[L_t = L]$ for all values of t from 0 to a certain time horizon T and for all possible levels of loss. We discuss how this assumption relates to the actual information available from the market in Section 3.2.

We now reduce Eq. (2) to a simpler forward Kolmogorov equation written on the density of loss, $P(L,t)$. This reduction is achieved by an integration of both parts of Eq. (2) over λ, which leads to

$$\frac{\partial P(L,t)}{\partial t} = \mathbf{1}_{L \geq h} \Lambda(L-h,t) P(L-h,t) - \Lambda(L,t) P(L,t). \tag{6}$$

The quantity $\Lambda(L,t)$, which we call the *local* intensity, is defined by the equation

$$\Lambda(L,t) P(L,t) = \int_0^\infty \lambda \, p(\lambda, L, t) d\lambda, \tag{7}$$

and has the meaning of the expectation of the intensity λ_t conditional on the loss L accumulated by the time t,

$$\Lambda(L,t) = \mathrm{E}[\lambda_t | L_t = L]. \tag{8}$$

We obtained Eq. (8) from a particular equation for the stochastic evolution of the intensity λ_t. It can be shown that the result is more general and is also valid for an adapted intensity process which is not necessarily Markovian. A more detailed discussion can be found in Appendix B.[3]

Eq. (6) is the forward Kolmogorov equation of the jump process based on the intensity $\Lambda(L_t, t)$ considered as a stochastic process. This jump process is known as a continuous-time, non-homogeneous, one-step Markov chain. The state space of this Markov chain is given by the grid of possible loss values, $0, h, 2h, \ldots, N_{\max} h$. For every $L < N_{\max} h$, the intensity of the transition $L \to L+h$ at time t is equal to $\Lambda(L,t)$, while the intensities of all other transitions are 0. Viewed as an intensity-based jump process, the one-step Markov chain is a specific case with the intensity being a deterministic function of time and loss. By analogy with local volatility models of exchange rates or interest rates (see Appendix B), it is natural to call this model the *local intensity* model.

The local intensity model has recently been considered by van der Voort (2005) who applied it to the pricing of forward starting CDOs. This model also appears in the works of Sidenius et al. (2005) and Schönbucher (2005) who use it as a part of their constructions of the dynamic

[3] See also the forthcoming work by Giesecke (2007) which contains a systematic exposition and further applications of the method of Markovian projection to basket credit modeling.

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

framework. We regard the local intensity model as a very useful tool for the calibration of models with richer dynamics, but which, by itself, is generally insufficient as a dynamic model (see, e.g., the numerical results for tranche options in Section 4.2).

The local intensity $\Lambda(L,t)$ can be easily calibrated to a given density of loss $P(L,t)$, which is why van der Voort called it an implied loss model. Indeed, summing up Eq. (6) from $L=0$ to $L=K=kh$ for any $k \le N_{\max}$, we obtain

$$\frac{\partial}{\partial t}\sum_{L=0}^{K} P(L,t) = -\Lambda(K,t)P(K,t), \tag{9}$$

which uniquely determines the local intensity $\Lambda(K,t)$ provided $P(K,t) \neq 0$ (that is, for all states which can be achieved by the process),

$$\Lambda(K,t) = -\frac{1}{P(K,t)}\frac{\partial \mathrm{P}[L_t \le K]}{\partial t}. \tag{10}$$

This completes the first step of the calibration procedure.

The next step is to find the function $\rho(L,t)$ consistent with the local intensity $\Lambda(L,t)$. In order to accomplish this task, we take the time derivative of Eq. (7),

$$\int_0^\infty \lambda \frac{\partial p(\lambda,L,t)}{\partial t} d\lambda \ = P(L,t)\frac{\partial \Lambda(L,t)}{\partial t} + \Lambda(L,t)\frac{\partial P(L,t)}{\partial t}, \tag{11}$$

and substitute the time derivatives of the densities $p(\lambda,L,t)$ and $P(L,t)$ from Eqs. (2) and (6) respectively. The resulting equation can be formally solved for $\rho(L,t)$ (again, for all accessible states which obey $P(L,t) \neq 0$), to give

$$\begin{aligned}
\rho(L,t) &= \Lambda(L,t) + \frac{1}{\kappa}\frac{\partial \Lambda(L,t)}{\partial t} \\
&+ \frac{\Lambda(L,t)\left(\mathbf{1}_{L\ge h}\Lambda(L-h,t)P(L-h,t) - \Lambda(L,t)P(L,t)\right)}{\kappa P(L,t)} \\
&+ \frac{M(L,t) - \mathbf{1}_{L\ge h}M(L-h,t)}{\kappa P(L,T)},
\end{aligned} \tag{12}$$

where

$$M(L,t) = \int_0^\infty \lambda^2 p(\lambda,L,t) d\lambda. \tag{13}$$

Eq. (12) is not an analytical solution for $\rho(L,t)$ because this function is implicitly present in the last term in the right-hand side via Eq. (2) that determines the density $p(\lambda,L,t)$. Nevertheless, Eq. (12) essentially solves the calibration problem. Indeed, a substitution of the function $\rho(L,t)$ from Eq. (12) into Eq. (2) leads to an integro-differential equation for the density $p(\lambda,L,t)$, which can be solved numerically. After that, the function $\rho(L,t)$ can be restored from Eq. (12).

Practically, instead of writing down the explicit integro-differential equation, it is more convenient to use Eqs. (12), (13), and (2) to set up a recursive procedure. We illustrate this procedure using a simple first order scheme. We discretize the time dimension into a sequence of small intervals, $[0,t_1]$, $[t_1,t_2]$,..., and introduce a discretization for λ (the loss variable being already discrete). Starting at $t=0$ with the density $p(\lambda,L,t)$ specified by the boundary

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

condition (4), we find the corresponding $M(L,0)$ from Eq. (13) and plug the result into Eq. (12) to obtain $\rho(L,0)$. Eq. (2) is then used to propagate the density to $t = t_1$, and then the entire procedure is repeated for each $t = t_i$ until the desired maturity of the model is reached. We note that only one complete propagation through the three-dimensional lattice of the discretized values t_{i_1}, λ_{i_2}, and $L_{i_3} = i_3 h$ is required. We also note that the integration over λ in Eq. (13) does not lead to any significant performance overhead. The overall computational effort is similar to that involved in the calibration of a typical two-factor interest rate model.

It is also possible to use higher order discretization schemes. In our numerical experiments, second order schemes performed best.

3.2 From market data to local intensity

We now turn to the calibration of the local intensity $\Lambda(L,t)$ to the actual market data, that is, to single tranche CDOs. (For a brief description of single tranche CDOs, see Appendix A.) We see two alternatives for the local intensity calibration.

In the approach by van der Voort (2005), the relationship (10) is used to find the local intensity from the probability distribution of loss. In turn, the probability distribution of loss is found assuming a particular factor model with a sufficient number of degrees of freedom to fit all the tranches as well as the underlying credit index. For example, this could be the random factor loadings model of Andersen and Sidenius (2005a), or a mixture of several Gaussian copulas considered by Li and Liang (2006), or the implied copula model of Hull and White (2006). Dynamical properties of the auxiliary factor model are ignored, the only requirement being that the model should produce an arbitrage-free distribution of loss. (This requirement generally disqualifies the use of the base correlations model.)

Alternatively, one can assume a certain functional form for the local intensity $\Lambda(L,t)$ and do a parametric fit to the tranches and the index. In our opinion, this approach is more suitable for the purpose of dynamic model building since all assumptions about the time dependence of the local intensity are kept explicit. Such assumptions cannot be avoided because liquid tranche quotes are only available for a handful of maturity points. Consequently, we need to look for the most natural way to interpolate the local intensity. At present, we do not see any reliable way to control the time dependence of the local intensity within the static factor models approach. Therefore, we prefer dealing directly with a functional form of the local intensity.

We now proceed to parametric modeling of the local intensity. In doing so, we found it convenient to express the local intensity as a function $\Lambda(N,t)$ of the number of defaults N and time t instead of loss L and time t. (With a deterministic LGD, this is an equivalent description because the number of defaults, N, is proportional to loss, $L = Nh$.) The challenge is to reconcile the behavior near $t = 0$, which as we will see turns out to be singular, with the more regular behavior away from $t = 0$. We address this issue by modeling the N-dependence of the local intensity as a power series in the following parameter,

$$x = \frac{N}{\bar{N}(t) + z(t)}. \tag{14}$$

The function $\bar{N}(t)$ is the average number of defaults occurred by time t, which is related to (but not fully fixed by) the quote for the credit index spread. The function $z(t)$ is introduced to regularize the singularity at $t \to 0$ in Eq. (14). Specifically, we used an exponential function,

$$z(t) = \exp(-\gamma t), \tag{15}$$

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

but any monotonic function that decays sufficiently fast with time could be used instead. A representation of the local intensity in terms of the parameter x ensures that, for $t \gg \gamma^{-1}$, the local intensity becomes a function of the number of defaults, N, normalized by the expected number of defaults, $\bar{N}(t)$. This normalization reflects the fact that the typical number of defaults naturally grows with time even in the case where the typical local intensity stays constant in time.

Thus, we look for the local intensity in the form

$$\Lambda(N,t) = \alpha_0(t) + \alpha(N,t), \qquad \alpha(N,t) = \sum_{p=1}^{p_{max}} \alpha_p(t)\, x^p. \tag{16}$$

To ensure the proper normalization of the probability density within the loss interval $[0, hN_{\rm max}]$ we assume that the above definition holds for $N < N_{\rm max}$ while at the boundary $N = N_{\rm max}$ the local intensity is zero

$$\Lambda(N_{\rm max}, t) = 0. \tag{17}$$

The main dependence of the local intensity on time in Eq.(16) is contained in the parameter x. A residual dependence on time in the coefficients α_p is included to ensure the matching with the initial condition at $t = 0$ (discussed below), and also for the fit of tranches with different maturities. The number of parameters $p_{\rm max}$ should be less or equal to the number of tranches at a single maturity, $N_{\rm tr}$. The best quality is usually achieved at $p_{\rm max} = N_{\rm tr}$. [4]

The choice of the initial condition for the local intensity $\Lambda(N, 0)$ has to be consistent with the initial distribution $p_0(\lambda)$ assumed in Eq. (4) for the stochastic intensity model. Indeed, solving the equation (2) near $t = 0$, we obtain a family of Poisson distributions with constant intensities λ distributed with the density $p_0(\lambda)$,

$$p\,(\lambda, N, t \approx 0) = \frac{(\lambda t)^N \exp(-\lambda t) p_0(\lambda)}{N!}. \tag{18}$$

In this equation, $\exp(-\lambda t)$ can be replaced by 1 in the same order of approximation. Using this expression for the density $p(\lambda, L, t)$ with Eqs. (5) and (7), we obtain the following initial condition for the local intensity

$$\Lambda(N, 0) = \frac{\int \lambda^{N+1} p_0(\lambda) d\lambda}{\int \lambda^N p_0(\lambda) d\lambda} \tag{19}$$

in terms of the moments of the initial stochastic intensity distribution.

Setting $N = 0$ in Eq. (19), we obtain a correct relationship for the instantaneous intensity of the first default,

$$\Lambda(0, 0) = \int \lambda\, p_0(\lambda) d\lambda. \tag{20}$$

The relevance of the relationship (19) with $N \geq 1$ is not immediately obvious because neither the local intensity $\Lambda(N, 0)$ nor the higher moments of the initial intensity distribution $p_0(\lambda)$ can be extracted from the market data. Nevertheless, Eq. (19) is very important for the consistency of the calibration scheme. Indeed, it shows that a particular choice for the initial distribution

[4] To avoid too big values of the local intensity arising from high powers of x in the expansion (16) we used the regularization $\alpha(N,t) \to \alpha(N,t)\, \alpha_{\rm max}/(\alpha(N,t) + \alpha_{\rm max})$, with $\alpha_{\rm max} \sim 100$.

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

$p_0(\lambda)$ fully determines the initial condition for the local intensity $\Lambda(N,0)$, and vice versa. We note also that the distribution $p_0(\lambda)$ is not guaranteed to exist for an arbitrary choice of $\Lambda(N,0)$; in particular, it is not possible to set $\Lambda(N,0)=0$ for $N\geq 1$, which might seem natural.

The easiest way to ensure that Eq. (19) holds is to choose a particular distribution $p_0(\lambda)$ and restore $\Lambda(N,0)$. We used the simplest choice corresponding to a deterministic initial condition $\lambda=\lambda_0$ for the stochastic intensity,

$$p_0(\lambda)=\delta(\lambda-\lambda_0). \tag{21}$$

(Here $\delta(x)$ is the Dirac δ-function, describing a unit weight localized at $x=0$). This corresponds to a constant initial condition for the local intensity,

$$\Lambda(N,0)=\lambda_0, \tag{22}$$

which leads to $\alpha_0(0)=\lambda_0$ and the following set of initial conditions for the coefficients α_p with $p\geq 1$ in Eq. (16),

$$\alpha_p(0)=0, \qquad p\geq 1. \tag{23}$$

These initial conditions are satisfied by the following time dependence,

$$\alpha_p(t)=\tilde{\alpha}_p\,(1-z(t)), \qquad p\geq 1, \tag{24}$$

which interpolates between 0 and an asymptotic constant value. The numerical values of the coefficients $\tilde{\alpha}_p$ can be fitted to a set of tranches with a single maturity. Simultaneous fit to tranches with several different maturities can be achieved using an additional term structure of the coefficients $\tilde{\alpha}_p$.

We finally show that the term $\alpha_0(t)$ is fixed by the time dependence of the average number of defaults $\bar{N}(t)$. Starting from the expression

$$\bar{N}(t)=\sum_{N>0}^{N_{max}} P(N,t)\,N, \tag{25}$$

we take the time derivative of both sides and use the Markov chain equation (6) to obtain

$$\frac{d\bar{N}(t)}{dt}=\sum_{N>0}^{N_{max}} N\left(\Lambda(N-1,t)P(N-1,t)-\Lambda(N,t)P(N,t)\right). \tag{26}$$

Substituting the expression (16) for $\Lambda(N,t)$, and using that $\Lambda(N_{max},t)=0$, we find

$$\alpha_0(t)=\frac{\frac{d\bar{N}(t)}{dt}-\sum\limits_{N\geq 0}^{N_{max}-1}\alpha(N,t)P(N,t)}{1-P(N_{max},t)}. \tag{27}$$

This equation is used to determine $\alpha_0(t)$ while solving the forward equation (6) for the local intensity model. We note also that the initial condition $\lambda=\lambda_0$ for the stochastic intensity is given by

$$\lambda_0=\alpha_0(0)=\frac{d\bar{N}(t)}{dt}\Big|_{t=0}. \tag{28}$$

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

Table 1 here

3.3 Numerical results

We present numerical results for the calibration to a set of tranches on Dow Jones CDX.NA.IG.7 5Y quoted on Jan 12, 2007 (see Table 1). The index was quoted at 33.5bp.

We first need to fit the average number of defaults $\bar{N}(t)$ to the index. To fix the time dependence $\bar{N}(t)$ completely we would need to know a full term structure of index quotes for all maturities until 5y. In reality, the most one could currently get is a quote for the index at 3y and quotes for CDS spreads for some of the assets in the basket at 1y and 3y.

In the absence of direct access to reliable information about the initial segment of the term structure, we are forced to introduce a parametric functional dependence for $\bar{N}(t)$. It turns out that a simple exponential decay of the fraction of survived assets, $1 - \bar{N}(t)/N_{\max}$, with a constant hazard rate does not allow for a robust calibration to the tranches. Therefore, we introduced a slightly more complicated form

$$\bar{N}(t)/N_{\max} = 1 - e^{-a(t+bt^2)}, \tag{29}$$

where a and b are fitting parameters. A positive coefficient b takes into account the effect of an upward slope of the spread curve. We tried different values of b, in each case solving for the value of a necessary to reproduce the spread of the index.

Once the dependence $\bar{N}(t)$ has been fixed, we fit the model to the tranches by adjusting the coefficients α'_p that determine the local intensity according to Eqs. (16) and (24). This is done using a multi-dimensional solver. For the calculation of tranche values, we used Eqs. (39) and (40) from Appendix A with $\alpha = \beta = 0.5$ and the standard assumption of 40% for the recovery rate. (The corresponding value for the LGD is $h = 0.6$.) The values of the stop-loss options in the local intensity model are obtained from the loss probability density, $P(L,t)$, found by integrating Eq. (6).

We observed that the quality of the fit was sensitive to the shape of the time dependence of the average number of defaults, $\bar{N}(t)$, controlled by the parameter b. In particular, we were not able to fit all five tranches with the required accuracy for $b = 0$. However, increasing b to the values of the order of 1.0 to 5.0 resulted in a dramatic improvement of the quality of the fit. For example, we were able to match all the spreads with the accuracy of 0.3 bp using $a = 0.00049$, $b = 4$ and four terms in the local intensity expansion series (16), as shown in Table 1. We used $\gamma = 1.6$ for the interpolation scale which is defined by Eq. (15) and enters Eqs. (14) and (24). The quality of the fit was essentially insensitive to this parameter.

Figure 1 here

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

The surface of the resulting local intensity $\Lambda(N,t)$ is plotted in Fig. 1. There is a spike in the region of small values of t and large values of N, which, however, does not lead to any serious numerical difficulties because the probability of reaching this region is vanishingly small. For a fixed value of t, the local intensity strongly increases with the number of defaults and has a positive convexity. It can be demonstrated that this shape is a signature of the skew in the Gaussian base correlations. For example, the local intensity surface derived from a Gaussian copula with constant correlations will have a much flatter shape (see Appendix B for an additional discussion).

Now that the local intensity surface is known, we can proceed to the final step in the calibration of the stochastic model and find the function ρ using the method described in Section 3.1. The resulting function ρ depends on the values of the parameters κ and σ in Eq. (1). In Fig. 2, we present a typical surface plot of the function $\rho(N,t)$, using $\kappa = 1$, $\sigma = 1$, and the number of defaults, N, instead of loss, L, as an argument. The qualitative behavior of $\rho(N,t)$ is similar to that of the local intensity $\Lambda(N,t)$, with a spike in the region of large N and small t, which is again irrelevant because of a negligible probability of reaching this region.

Figure 2 here

---

4 Dynamic applications

We now turn to the pricing of dynamics-sensitive financial instruments with the stochastic model defined by Eq. (1). An efficient implementation is possible both for the forward simulation and for the backward induction, the latter because the model is low-dimensional Markovian. In the present work, we focus our attention on the evaluation of tranche options using backward induction. Applications to forward starting CDOs and other instruments that require a forward simulation are deferred to a separate work.

4.1 Backward induction

We begin with a generic description of backward induction, assuming that the discounting rate is 0, so that all discount factors are equal to 1. (We will restore the proper discounting in Section 4.2.) Let $F(\lambda, L, T)$ be an arbitrary payoff function of the pair of state variables (λ, L) that can be achieved at time T. The backward induction to an earlier time t is the procedure of going from $F(\lambda, L, T)$ to another payoff function $F(\lambda, L, t)$ defined as a conditional expectation with respect to the state achieved at time t,

$$F(\lambda, L, t) = \mathrm{E}\left[F(\lambda_T, L_T, T) | L_t = L, \lambda_t = \lambda\right]. \tag{30}$$

This expectation satisfies the backward Kolmogorov equation

$$\frac{\partial F(\lambda, L, t)}{\partial t} = -\hat{A}_{back}\, F(\lambda, L, t), \tag{31}$$

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

where the action of the generator $\hat{A}_{back}$ on an arbitrary function $F(\lambda, L, t)$ is defined by

$$\hat{A}_{back}\, F(\lambda, L, t) = \left(\kappa(\rho(L,t) - \lambda) \frac{\partial}{\partial \lambda} + \frac{1}{2}\sigma^2 \lambda \frac{\partial^2}{\partial \lambda^2} \right) F(\lambda, L, t) + \lambda \left(F(\lambda, L + h, t) - F(\lambda, L, t) \right). \tag{32}$$

This generator is a conjugate of the generator present in the right-hand side of the forward Kolmogorov equation (2). Correspondingly, our numerical solution of the discretized backward Kolmogorov equation is a conjugated version of the solution of the forward Kolmogorov equation outlined in Appendix C.

It follows from the replication arguments presented in Appendix A that the payoff fundamental for the tranche valuation is that of the stop-loss option. For a stop-loss option with maturity T and strike X, the payoff is a deterministic function of state at time T,

$$P_{T,X}(\lambda, L, T) = (L_T - X)^+. \tag{33}$$

There is no dependence on λ in the right-hand side of Eq. (33). Such dependence will appear after a backward induction to an earlier time t, the result of which represents the value of the stop-loss option as viewed from the prospective of the time t,

$$P_{T,X}(\lambda, L, t) = \mathrm{E}[(L_T - X)^+ | L_t = L, \lambda_t = \lambda]. \tag{34}$$

Taking t to be the exercise time, the value of the entire tranche as of this time can be represented as a linear combination of the quantities (34) with different values of X and T (see Appendix A). In order to evaluate, for example, an option to enter the tranche, we only need to take the positive part and perform a final backward induction to $t = 0$.

4.2 Numerical results for the tranche option

We consider an option that gives the right to buy the protection leg of the tranche, selling the fee leg with a fixed value of spread, called strike, on a certain exercise date T_{ex}. As discussed above, the payoff from the exercise of the option can be represented as a function $V(\lambda, L, T_{ex})$ of state achieved at time T_{ex}. More specifically, the payoff is given by a linear combination of the elementary conditional expectations (34) for a portfolio of stop-loss options,

$$V(\lambda, L, T_{ex}) = \sum_{T_i \geq T_{ex}} w_i P_{T_i, X_i}(\lambda, L, T_{ex}) D(T_i)/D(T_{ex}). \tag{35}$$

The weights w_i are defined in Appendix A. Non-trivial discounting factors $D(t)$ have been restored under the assumption of deterministic interest rates. (Extending the model to include additional dimensions for stochastic interest rates is a subject for future work.)

The option exercise condition is taken into account by replacing negative values of $V(\lambda, L, T_{ex})$ by zero,

$$V^+(\lambda, L, T_{ex}) = \max\left(V(\lambda, L, T_{ex}), 0\right). \tag{36}$$

The value of the option is finally obtained by applying the backward induction from T_{ex} to 0 to $V^+(\lambda, L, T_{ex})$ and multiplying by the discount factor $D(T_{ex})$.

We note that the same backward induction procedure can also be implemented for the local intensity model. The only difference is that the states of the local intensity model include loss

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

L but not intensity λ. The local intensity model can also be regarded as a limit of the two-dimensional model at $\sigma \to 0$, $\kappa \to \infty$. We will give option pricing results produced by this model to compare with those obtained within the full model.

Figure 3 here

The dependence of the option value on the strike spread is shown in Fig. 3 for the case of a mezzanine tranche 3-7% with 159 days to exercise. The at-the-money (ATM) strike is defined as the model-independent forward spread S_F, which can be obtained by dividing the forward value of the protection leg by the basis point value of the fee leg. In our case, $S_F = 79.7$ bp. Solid curves correspond to $\kappa = 1$ and different values of the parameter $D = \sigma^2/2$. The value of the option in the local intensity model is shown by the dashed line. One can see that the change of the strength of the diffusion term in the stochastic intensity model leads to a noticeable change in option values, thereby providing some freedom to fit market quotes for the options. This is in contrast to the local intensity model which does not have any free parameters remaining after the calibration to the tranches.

Figure 4 here

Fig. 4 provides an equivalent representation of the option prices in terms of implied Black volatilities. The order of magnitude 80%-120% for the volatilities in the ATM region is consistent with typical values used in heuristic estimates.

The hockey-stick-like dependence of the option price generated by the local intensity model is in agreement with the general intuition about the zero-volatility limit. We note, however, that the local intensity model retains stochastic degrees of freedom of the jump process even though it can be obtained as a degenerate limit of the full two-dimensional model. The appearance of two straight lines can be understood by taking into account that the probability of a default before the exercise date is low, so that the main contribution to the option price comes from the scenarios with either 0 or 1 default before T_{ex}. In each of the two scenarios, the option is either worthless or depends linearly on the strike. The initial steep segment comes from the no-default scenario. The nearly flat segment corresponds to the scenario where the strike is large but the option remains valuable because of the default before exercise. (We assume here that the option does not knock-out on default.)

We can conclude that the local intensity model is too inflexible to provide a good description of tranche options. The complete two-dimensional model does not suffer from the type of degeneracy exhibited by the local intensity model because of the smoothing provided by an integration over a continuous range of local intensities λ.

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

It is interesting to note that the value of the option decreases with increasing parameter $D = \sigma^2/2$. This behavior is not intuitive as the value of the option usually grows with volatility. One should keep in mind, however, that the stochastic model is calibrated to remain consistent with the same surface of loss for any value of σ. An increased volatility of the diffusion term is compensated by a decrease in the strength of the back action term driven by the function $\rho(L, t)$. The direction of the total effect on the option value is not obvious.

5 Conclusions

We suggested a new intensity-based Markovian model for the dynamics of the aggregate credit loss and developed an efficient method for the calibration to the distribution of loss implied by the market data for CDO tranches. The calibration method is based on the technique of Markovian projection, which in our case allows us to associate the original two-dimensional model with a Markov chain generated by a local surface of default intensity. The Markov chain model is used on the first step of the calibration procedure to find the local intensity and the distribution of loss consistent with the market spreads of CDO tranches. After that, the full two-dimensional stochastic model is calibrated to the local intensity, which already includes all the necessary market information.

Apart from the ability to match a generic distribution of loss, our model has additional parametric freedom to control the fit to more complicated dynamics-sensitive instruments. Specifically, the parameter σ controls the strength of diffusive fluctuations of default intensity while the parameter κ sets the time scale of reversion in the drift term. The SDE for the intensity is similar to that for the short rate in the CIR model. The similarity, however, should be explored with caution because the drift term includes a back action of the loss process onto the intensity process. Changing the relative value of the coefficients κ and σ, we can go from an intensity process dominated by diffusion to one dominated by the back action of loss, while maintaining the calibration to the same distribution of loss.

The model can be used for pricing of different financial instruments via standard methods developed for Markovian stochastic processes. In the present paper, we focused on applications to tranche options. This instrument can be conveniently evaluated using the backward induction technique. We found that the model can produce a wide range of option prices corresponding to different values of σ.

We note that our approach is not limited to the specific CIR-like intensity dynamics (1). Other equations, for example those based on BK-like evolution, may turn out to be more suitable for the purpose of credit portfolio modeling. The current evidence, however, indicates that a sufficiently flexible form of a back action term is essential for a model's ability to match the market of CDO tranches in a robust way.

Acknowledgements

We are grateful to the organizers and the participants of the 5th Annual *Advances in Econometrics* conference at LSU, Baton Rouge (November 3–5, 2006) for the opportunity to present and discuss the results of this work. We would like to thank Alexandre Antonov, Leonid Ryzhik, Serguei Mechkov, Ren-Jie Zhang for useful discussions and especially Kay Giesecke for his valuable comments on the generalization of Markovian projection technique to the case of processes

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

which are not doubly stochastic. We are grateful to Gregory Whitten and our colleagues at NumeriX for support of our work.

Appendices

A Single tranche CDO

The purpose of this section is to introduce single tranche CDOs and justify the replication of a tranche by a portfolio of stop-loss options. A single tranche CDO is a synthetic basket credit instrument which involves two parties and references a portfolio of credit names. One party is the buyer of the protection, the other is the seller of the protection. A single tranche CDO contract defines two bounds, $k < K$, called attachment points and usually quoted as percentage points of the total original reference notional A of the underlying portfolio. The lowest tranche, $0-3\%$ or similar, is customarily called the equity tranche. The highest tranche, $30\%-100\%$ or alike, is called supersenior. The other tranches are called mezzanine and senior. The difference of the bounds $K-k$ is the original notional of the tranche, which is the cap of the liability held by the seller of the protection. Additionally, the single tranche CDO contract defines a schedule of accrual and payment dates, a fixed annualized periodic rate S, called tranche spread, and, in the case of equity tranches, an upfront payment to be made by the buyer of the protection. Par spread of a single tranche CDO is defined as the value of S that makes the present value of the tranche equal to zero.

The cashflows are driven by the slice of the loss of the reference portfolio within the segment $[k, K]$. The total loss sustained by the tranche between its inception at time 0 and time T is given by a difference of two stop-loss options,

$$L_{k,K}(T) = (L(T) - k)^+ - (L(T) - K)^+ \tag{37}$$

(by definition $(x)^+ = x$ if $x > 0$ and $(x)^+ = 0$ otherwise). As soon as a positive jump $\Delta L_{k,K}$ in the quantity $L_{k,K}$ is reported, the seller of the protection must pay the amount $\Delta L_{k,K}$ to the buyer of the protection. This is the only source of the payments made by the seller of the protection. The payments made by the buyer of the protection are determined by the outstanding notional of the tranche $A_{k,K}(T)$ as a function of time T. The initial notional of the tranche is $A_{k,K}(0) = K - k$. The notional of the tranche at time T is given by

$$A_{k,K}(T) = A_{k,K}(0) - L_{k,K}(T). \tag{38}$$

The outstanding notional of the tranche is monitored every day of each payment period, and the fee is accrued on the outstanding notional of the tranche with the rate equal to the tranche spread S. The total accrued fee is paid by the buyer of the protection on the payment date.

Let the payment periods be $[0, T_1]$, $[T_1, T_2]$,$\dots$ $[T_{f-1}, T_f]$. Introducing the risk-free discount curve $D(t)$, the leg of the payments made by the protection seller (protection leg) can be approximated as

$$P_{\text{prot}} = \sum_i \left(\mathrm{E}[L_{k,K}(T_i)] - \mathrm{E}[L_{k,K}(T_{i-1})]\right) D(T_i). \tag{39}$$

Here we ignored the exact timings of defaults. This approximation can be refined by introducing a schedule of default observations, which is more frequent than the payment schedule.

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

The leg of the payments made by the protection buyer (fee leg) can be approximated as

$$P_{\text{fee}} = \sum_i S \cdot \tau(T_{i-1}, T_i)(\alpha A_{k,K}(T_{i-1}) + \beta A_{k,K}(T_i))D(T_i). \tag{40}$$

Here $\tau(T_{i-1}, T_i)$ is the daycount fraction from T_{i-1} to T_i, and α, $\beta = 1 - \alpha$ are the weights introduced to take into account the timing of defaults. When we set $\alpha = \beta = 0.5$ we effectively assume that the defaults on the average happen in the middle of the payment period. Again, it is possible to use a more frequent grid of observations to improve the accuracy of the calculation.

The present value of the tranche, P_{tr}, is equal to the difference of the legs, that is, $P_{\text{prot}} - P_{\text{fee}}$ for protection buyer, and $P_{\text{fee}} - P_{\text{prot}}$ for protection seller. It is easy to see that the final expression can be represented as a linear combination of stop-loss expectations,

$$P_{\text{tr}} = \sum_j w_j \mathrm{E}[(L_{t_j} - X_j)^+)]D(t_j). \tag{41}$$

Here t_j is either one of the payment dates or one of the dates of a more frequent grid introduced to improve the accuracy of the calculation; the strike X_j is one of the two attachment points, k or K, and w_j is a weight that can be positive or negative. We assume that the interest rates are deterministic and, where necessary, include the ratios of the discount factors into the definition of the weights w_j to obtain the replication in form (41).

The formula (41) is given in terms of unconditional expectations and, strictly speaking, does not express the static replication which has to hold at every moment in the life of the instrument. However, exactly the same derivation can be repeated with conditional expectations, leading to a static replication of the tranche by a portfolio of short and long positions in stop-loss options with the weights w_j.

In the case of super-senior tranches, it is also necessary to take into account the amortization provision. The obligatory amortization begins as soon as the cumulative recovery amount $R(T)$ exceeds $A-K$ and can extend to the total original notional of the tranche, $K-k$. The reduction of the tranche notional due to amortization is given by

$$R_{k,K}(T) = (R(T) - (A - K))^+ - (R(T) - (A - k))^+. \tag{42}$$

This quantity should be subtracted from the right-hand side of Eq. (38). It follows that a static replication of super-senior tranches requires recovery options in addition to stop-loss options. This complication does not limit the applicability of static factor models because the recovery options are also insensitive to the dynamics. We note, furthermore, that in the case of a deterministic LGD, the ratio of recovery and loss is a constant factor, so that the recovery options can be rewritten in terms of stop-loss options, which removes the need to model the process for the recovery separately.

We finally note that the index itself can be treated as a tranche with attachment points 0 and 100%. As with any tranche with a large value of the upper attachment point, it is necessary to take into account the contribution from recovery. The value of the index is fully determined by the term structure of expected loss, $\mathrm{E}[L(T)]$, and expected recovery, $\mathrm{E}[R(T)]$. Under the assumption of a deterministic LGD, the value of the tranche is fully determined by the term structure of the expected number of defaults $\bar{N}(t)$ (see Eq. (25)).

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

B Local volatility and local intensity

Here we discuss in more details the technique of Markovian projection for jump processes and establish a relationship between the stochastic intensity and the local intensity. We also elaborate on the analogy between the local intensity and the local volatility.

A stochastic volatility model involves a filtered probability space $(\Omega, \mathrm{P}, \{\mathcal{F}_t\})$ and an equation

$$dX_t = \alpha_t dt + \beta_t dW_t, \tag{43}$$

where the drift α_t and the volatility β_t are random processes adapted to the filtration $\{\mathcal{F}_t\}$. In a local volatility model, the processes α_t and β_t are deterministic functions of X_t and t. The local volatility model was introduced by Dupire (1994) in the form

$$dX_t/X_t = \mu(X_t, t)dt + \sigma(X_t, t)dW_t. \tag{44}$$

Local volatility models are regarded as a degenerate case of stochastic volatility models. They are easier to solve and calibrate to European options, but they do not generate very realistic dynamics.

There is a remarkable reduction of a stochastic volatility model to a local volatility model with the preservation of all one-dimensional marginal distributions, due to Gyöngy [14]. A non-technical statement of the claim is that the process X_t defined by Eq. (43) has the same one-dimensional distributions as the process Y_t defined by the equation

$$dY_t = a(Y_t, t)dt + b(Y_t, t)dW_t, \tag{45}$$

where $Y_0 = X_0$, and the local coefficients a, b are given by

$$a(x, t) = \mathrm{E}[\alpha_t | X_t = x], \tag{46}$$
$$b^2(x, t) = \mathrm{E}[\beta_t^2 | X_t = x]. \tag{47}$$

The mapping from the process X_t to the process Y_t is called the Markovian projection.

Piterbarg (2006) gave an intuitive proof and numerous applications of this result to various problems of mathematical finance. In such applications, the Markovian projection is typically used for fast calculation of European options in a stochastic volatility model. Fast calculation of European options is often critical to ensure adequate performance at the stage of model calibration. The method works because the European options only depend on the one-dimensional marginals of the underlying rate and can be computed in the effective local volatility model.

To extend the methodology of Markovian projection to stochastic intensity models of credit basket loss, we need a counterpart of Gyöngy's theorem for jump processes. Omitting the technical conditions, the statement is that a counting process N_t with an adapted stochastic intensity λ_t and $N_0 = 0$ has the same one-dimensional marginal distributions as the process M_t with the intensity $\Lambda(M, t)$ given by

$$\Lambda(M, t) = \mathrm{E}[\lambda_t | N_t = M]. \tag{48}$$

This is the same as Eq. (8) which was derived from the Kolmogorov equation for a specific stochastic intensity process. For a general proof,[5] we start with the expression (10) for the local

[5]The proof given in the first revision of our paper was applicable only to non self-affecting doubly stochastic processes. We thank Kay Giesecke for pointing this out.

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

intensity in terms of the probability distribution $P(N,t) = \mathrm{P}[N_t = N]$,

$$\Lambda(M,t) = -\frac{\partial \mathrm{P}[N_t \le M]/\partial t}{\mathrm{P}[N_t = M]}, \tag{49}$$

and write the derivative term as

$$\frac{d}{dt}\,\mathrm{P}[N_t \le M] = \frac{d}{dt}\mathrm{E}[\mathbf{1}_{N_t \le M}] = \lim_{\epsilon \to +0} \frac{1}{\epsilon}\mathrm{E}[\mathbf{1}_{N_{t+\epsilon} \le M} - \mathbf{1}_{N_t \le M}]. \tag{50}$$

Denote $\delta N = N_{t+\epsilon} - N_t$. Since $\delta N \ge 0$, the expression under the expectation in Eq. (50) can be written as

$$\mathbf{1}_{N_{t+\epsilon} \le M} - \mathbf{1}_{N_t \le M} = -\mathbf{1}_{M - \delta N < N_t \le M}. \tag{51}$$

We obtain

$$\frac{d}{dt}\,\mathrm{P}[N_t \le M] = -\lim_{\epsilon \to +0} \frac{1}{\epsilon}\mathrm{E}[\mathbf{1}_{M - \delta N < N_t \le M}]. \tag{52}$$

The leading contribution in ϵ comes from the realizations with $\delta N = 0, 1$. Thus, one can write

$$\begin{aligned} \frac{d}{dt}\,\mathrm{P}[N_t \le M] &= -\lim_{\epsilon \to +0} \frac{1}{\epsilon}\mathrm{E}[\delta N\, \mathbf{1}_{N_t = M}] \\ &= -\lim_{\epsilon \to +0} \frac{1}{\epsilon}\mathrm{E}[\delta N | N_t = M]\,\mathrm{P}[N_t = M] = -\Lambda(M,t)\mathrm{P}[N_t = M], \end{aligned} \tag{53}$$

which leads to Eq. (48).

We use the local intensity as a key element in the calibration procedure for the two-dimensional Markovian model. In concluding this section, we note that the local intensity calibrated to the market bears a distinctive signature of the correlation skew, as shown in Fig. 5. The Gaussian copula with any constant correlation value leads to a nearly linear dependence of the local intensity on the number of defaults. This is in contrast with the behavior of the local intensity calibrated to the actual market data, which shows a convex segment before saturating at a very large number of defaults.

Figure 5 here

C Discretization of intensity

Numerical integration of Eq. (2) by means of a finite difference scheme requires discretization of time t and intensity λ. The discretization of time does not pose any conceptual difficulties. The discretization of λ is more subtle because it needs to be done in a way preserving the key ingredients of the calibration method presented in Section 3.1, including Eqs. (6) and (12). Here we present a simple scheme that satisfies this requirement.

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

We use a uniform grid, $\lambda_i = i\,\Delta$, and introduce the finite difference operators $\hat{D}_\pm$ as

$$\hat{D}_+ f(\lambda_i) = \frac{f(\lambda_i + \Delta) - f(\lambda_i)}{\Delta}, \qquad \hat{D}_- f(\lambda_i) = \frac{f(\lambda_i) - f(\lambda_i - \Delta)}{\Delta}. \tag{54}$$

In the limit $\Delta \to 0$, these converge to the continuous derivative operator $d/d\lambda$. The discrete counterpart of the second order derivative $d^2/d\lambda^2$ reads

$$\hat{D}^2 = \hat{D}_+ \hat{D}_- = \hat{D}_- \hat{D}_+. \tag{55}$$

The discretized forward Kolmogorov equation (2) takes the form

$$\begin{aligned}\frac{\partial}{\partial t} p(\lambda_i, L, t) = \left(-\kappa \hat{D}_- (\rho(L,t) - \lambda_i) + \frac{1}{2}\sigma^2 \hat{D}^2 \lambda_i \right) p(\lambda_i, L, t) \\ + \lambda_i \left(p(\lambda_i, L - h, t) - p(\lambda_i, L, h) \right).\end{aligned} \tag{56}$$

(Here and below we omit the indicator $\mathbf{1}_{L \geq h}$ and assume that $p(\lambda_i, -h, t) = 0$.) Note that the term containing the first order derivative with respect to λ in Eq. (2) can be replaced either with $\hat{D}_+$ or with $\hat{D}_-$. With the choice of $\hat{D}_-$, we avoid the appearance of boundary terms after the summation over λ (see below). It is convenient to append $\lambda_{-1} = -\Delta$ to the range of allowed intensity values so that the boundary condition can be set as

$$p(\lambda_{-1}, L, t) = 0. \tag{57}$$

The probability density of loss and the local intensity in the discrete setting are defined similarly to Eqs. (5) and (7),

$$P(L,t) = \sum_{i=0}^{i_{\max}} p(\lambda_i, L, t), \tag{58}$$

$$\Lambda(L,t)\, P(L,t) = \sum_{i=0}^{i_{\max}} \lambda_i p(\lambda_i, L, t). \tag{59}$$

Summing both sides of Eq. (56) over λ_i from zero to the chosen limit $i_{\max}$, the forward Kolmogorov equation (6) is recovered. The boundary terms at the lower limit of summation disappear because of the condition (57).

We now proceed to the derivation of Eq. (12) in the discrete setting. Taking the derivative of both sides of Eq. (59) with respect to time, we get

$$\sum_{i=0}^{i_{\max}} \lambda_i \frac{\partial}{\partial t} p(\lambda_i, L, t) = P(L,t) \frac{\partial}{\partial t} \Lambda(L,t) + \Lambda(L,t) \frac{\partial}{\partial t} P(L,t). \tag{60}$$

After that, we insert the time derivatives of the distribution functions $p(\lambda_i, L, t)$ and $P(L,t)$ from Eqs. (6) and (56) into Eq. (60). We recover Eq. (12) using Eqs. (58) and (59) and the definition for the second moment of intensity,

$$M(L,t) = \sum_{i=0}^{i_{\max}} \lambda_i^2 \, p(\lambda_i, L, t). \tag{61}$$

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

Equation (56) represents a system of coupled ordinary differential equations that can be solved by any suitable method. We used the second order Runge-Kutta scheme.

We note that the absences of the probability flow to the region with negative intensities in the presented numerical scheme is guaranteed by construction. Thus, occasional appearances of negative values of $\rho(L,t)$ will not break the calibration algorithm.

The choice of the step, Δ, and the upper limit for the intensity, $i_{\max}\Delta$, is dictated by accuracy requirements. In our numerical experiments, we achieved the error of less than 1e-5 using $\Delta \sim 0.07 - 0.25$ and $i_{\max} \sim 1000$.

References

[1] Andersen, L. and J. Sidenius, 2005a, Extensions to the Gaussian copula: random recovery and random factor loadings. Journal of Credit Risk 1(1), 29–82.

[2] Andersen, L. and J. Sidenius, 2005b, CDO pricing with factor models: survey and comments. Journal of Credit Risk 1(3), 71–88.

[3] Andersen, L., 2006, Portfolio losses in factor models: Term structures and intertemporal loss dependence. Working paper, available at defaultrisk.com.

[4] Arnsdorf, M. and I. Halperin, 2007, BSLP: Markovian bivariate spread-loss model for portfolio credit derivatives. Working paper, available at defaultrisk.com.

[5] Bennani, N., 2005, The forward loss model: a dynamic term structure approach for the pricing of portfolio credit derivatives. Working paper, available at defaultrisk.com.

[6] Brigo, D., A. Pallavicini, and R. Torresetti, 2006, Calibration of CDO tranches with the dynamical generalized-Poisson loss model. Working paper, available at defaultrisk.com.

[7] Chapovsky, A., A. Rennie, and P. A. C. Tavares, 2006, Stochastic intensity modelling for structured credit exotics. Working paper, available at defaultrisk.com.

[8] Ding, X., K. Giesecke, and P. Tomecek, 2006, Time-changed birth processes and multi-name credit. Working paper, available at defaultrisk.com.

[9] Duffie, D., 2005, Credit risk modeling with affine processes. Journal of Banking and Finance 29, 2751–2802.

[10] Dupire, B., 1994, Pricing with a smile. Risk, January 1994, 18–20.

[11] Errais, E., K. Giesecke, and L. Goldberg, 2006, Pricing credit from the top down with affine point processes. Working paper, available at defaultrisk.com.

[12] Giesecke, K. and L. Goldberg, 2005, A top down approach to multi-name credit. Working paper, available at defaultrisk.com.

[13] Giesecke, K., 2007, From tranche spreads to default processes. Working paper.

[14] Gyöngy, I., 1986, Mimicking the one-dimensional marginal distributions of processes having an Itô differential. Probability Theory and Related Fields 71, 501–516.

[15] Hull, J. and A. White, 2006, Valuing credit derivatives using an implied copula approach. Journal of Derivatives 14 (2), 8–28.

[16] Li, D. and M. Liang 2006 CDO^2 pricing using Gaussian mixture model with transformation of loss distribution. Working paper, available at www.defaultrisk.com.

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

[17] Longstaff, F. A. and A. Rajan, 2006, An empirical analysis of the pricing of collateralized debt obligations. Working paper, available at www.defaultrisk.com.

[18] Piterbarg, V., 2006, Markovian projection method for volatility calibration. Working paper, available at SSRN: http://ssrn.com/abstract=906473.

[19] Schönbucher, P., 2005, Portfolio losses and the term structure of loss transition rates: a new methodology for the pricing of portfolio credit derivatives. Working paper, available at defaultrisk.com.

[20] Sidenius J., V. Piterbarg, and L. Andersen, 2005, A new framework for dynamic credit portfolio loss modelling. Working paper, available at defaultrisk.com.

[21] van der Voort, M., 2006, An implied loss model. Working paper, available at defaultrisk.com.

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

Spreads, bp		
Tranche	Model	Market
0-3%	500.2	500
3-7%	71.8	71.8
7-10%	13.3	13.3
10-15%	5.3	5.3
15-30%	2.8	2.6

Table 1: Market data and model calibration results for the spreads of the tranches on Dow Jones CDX.NA.IG.7 5Y quoted on Jan 12, 2007. The spread of the equity (0-3%) tranche assumes an upfront payment of 23.03% of the tranche notional. The coefficients in Eq. (24) found from the fit are $\tilde{\alpha}_1 = 0.06677$, $\tilde{\alpha}_2 = 0.07201$, $\tilde{\alpha}_3 = -0.006388$, $\tilde{\alpha}_4 = 0.00024$.

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

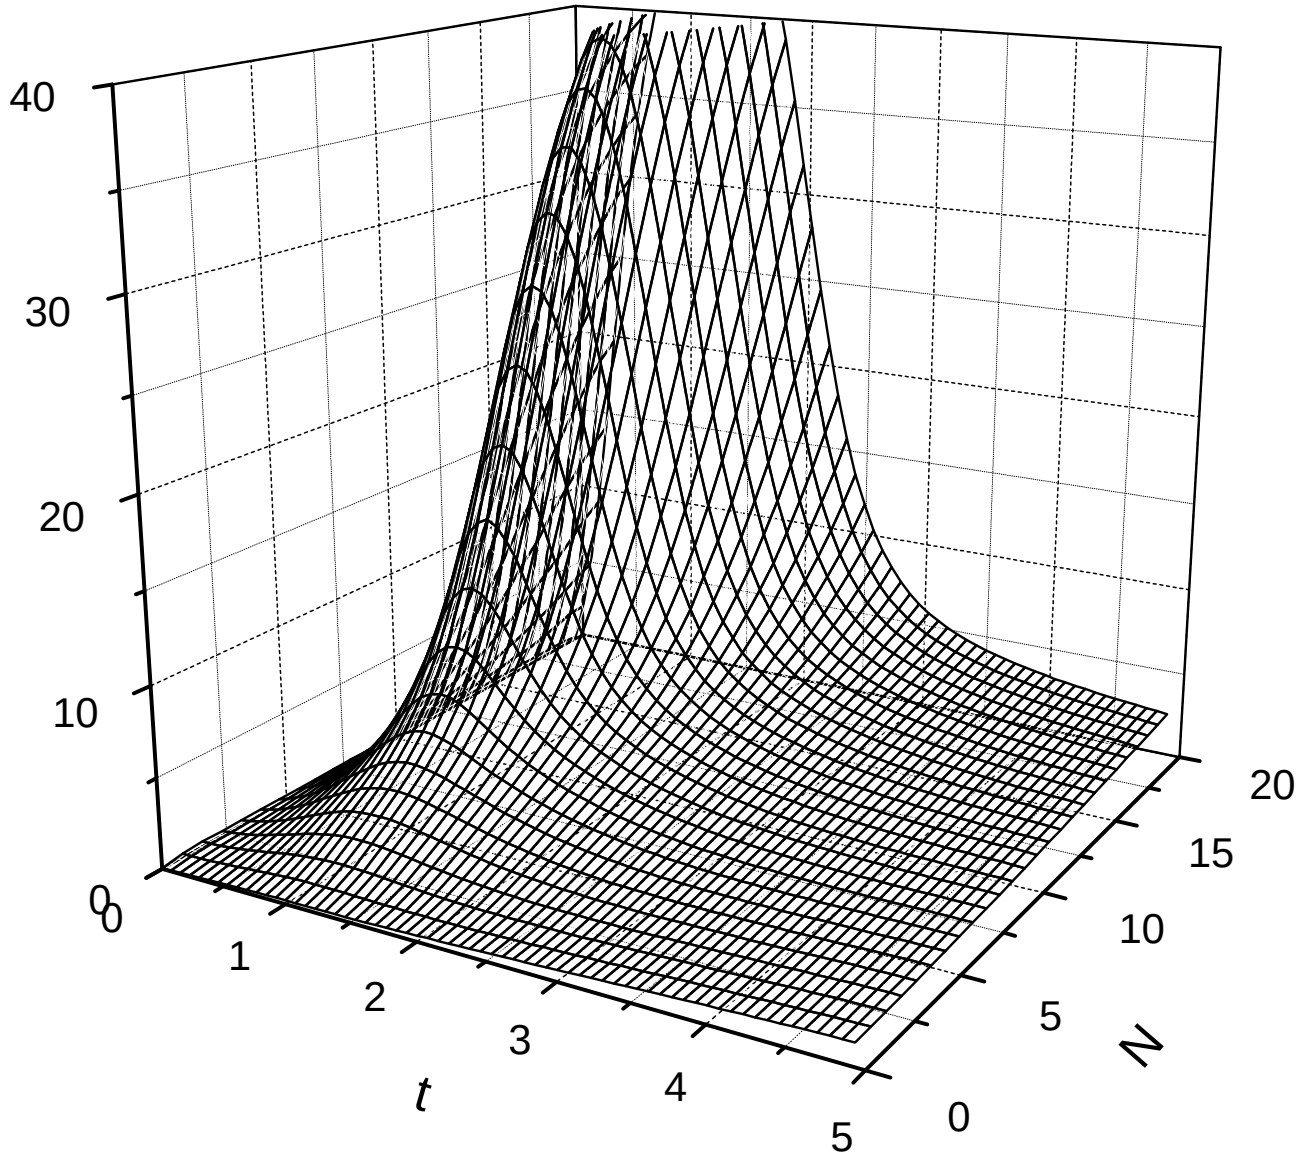

Figure 1: Local intensity as a function of the number of defaults N and time t measured in years, calibrated to the data of Table 1.

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

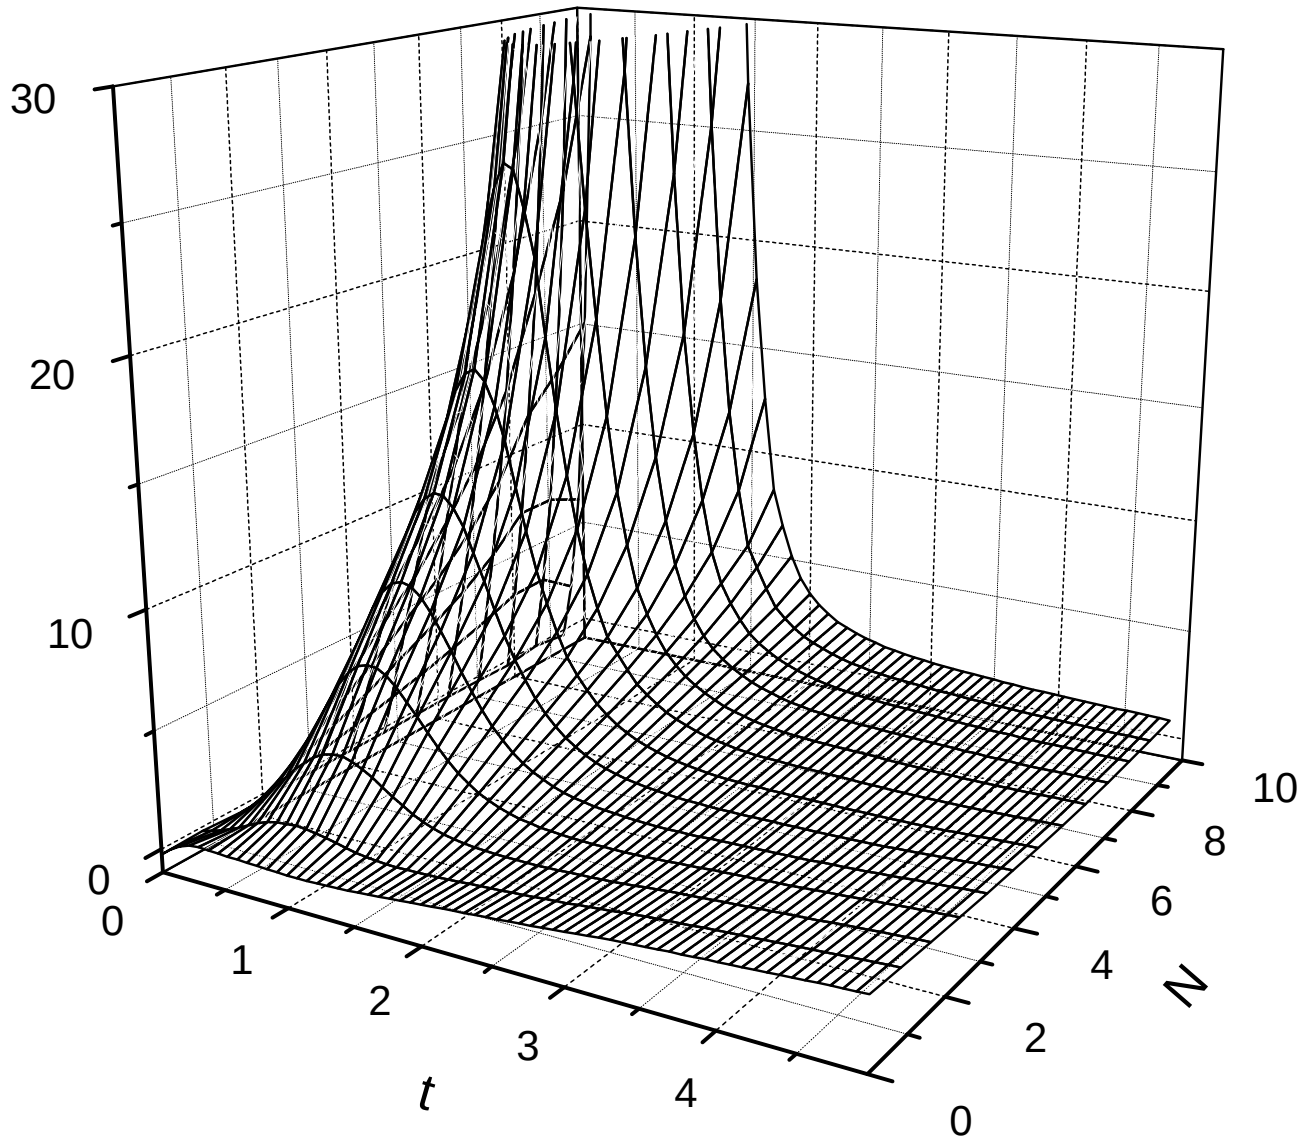

Figure 2: Dependence of the function ρ on the number of defaults N and time t measured in years for $\kappa = 1$ and $\sigma = 1$.

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

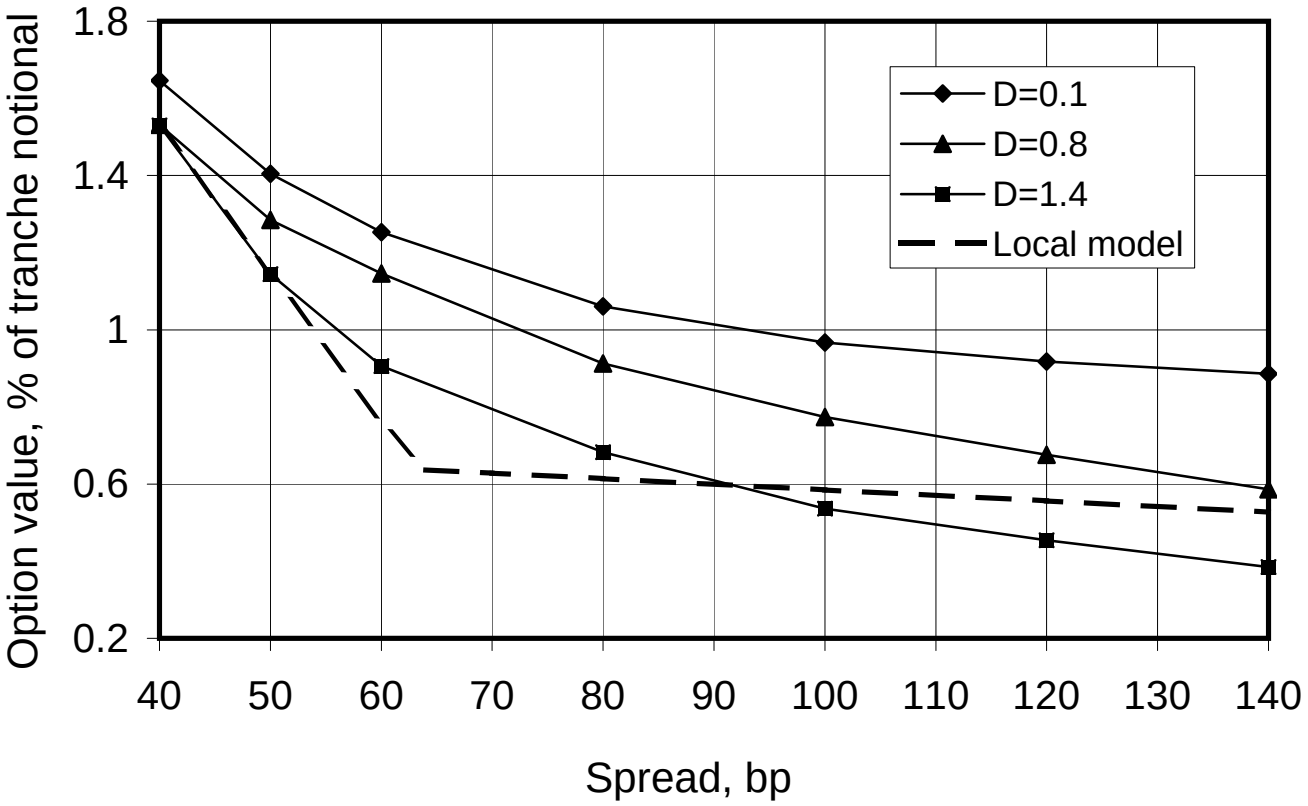

Figure 3: Dependence of the value of the option on the mezzanine 3-7% tranche on the strike spread. The time to exercise is 159 days, which corresponds to the exercise on June 20, 2007. Solid lines represent the results from the two-dimensional stochastic model with $\kappa = 1$ and $D = \sigma^2/2 = 0.1, 0.8, 1.4$. Dashed line is the result from the local intensity model. Option value is measured as a percentage of the tranche notional.

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

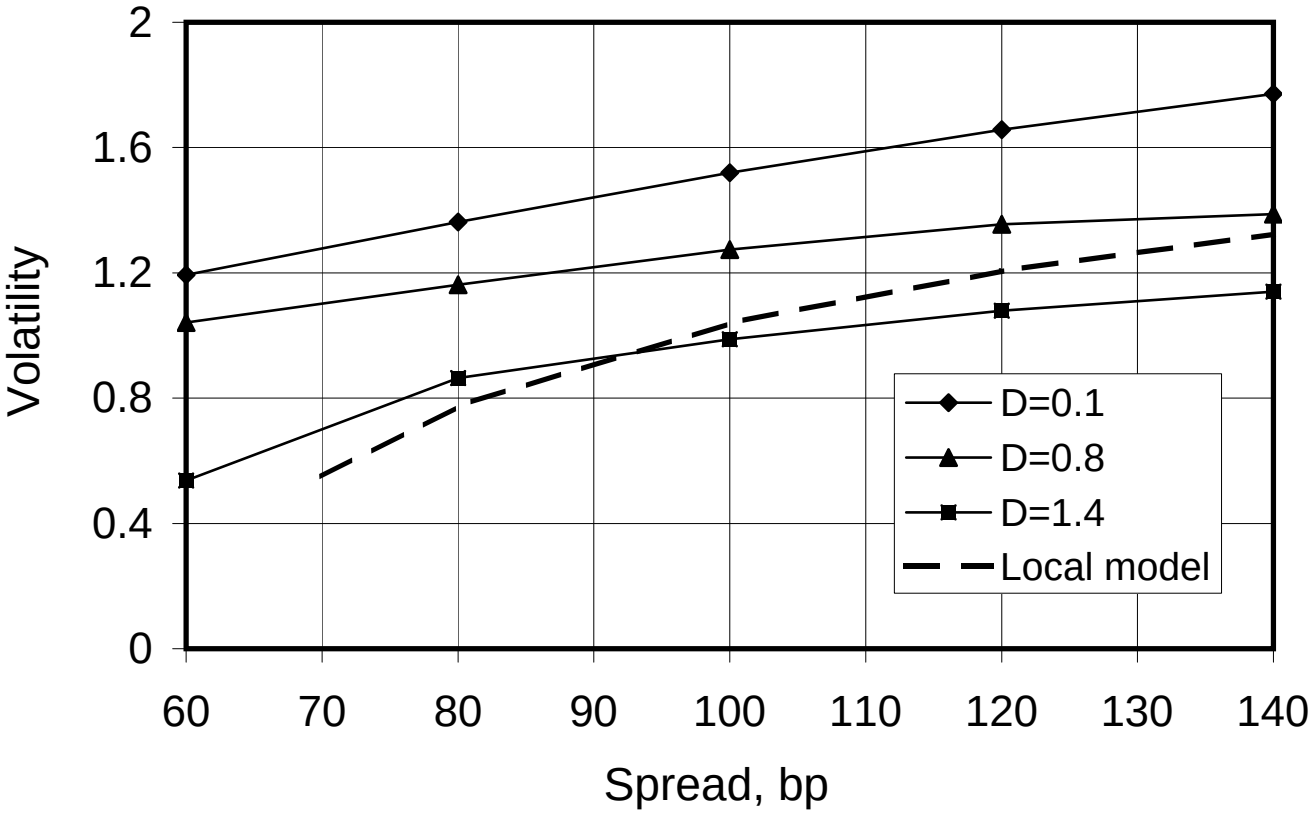

Figure 4: Value of the option on the mezzanine 3-7% tranche expressed in terms of implied Black volatilities. All option and model parameters are the same as in Fig. 3. Dashed line corresponds to the local intensity model.

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.

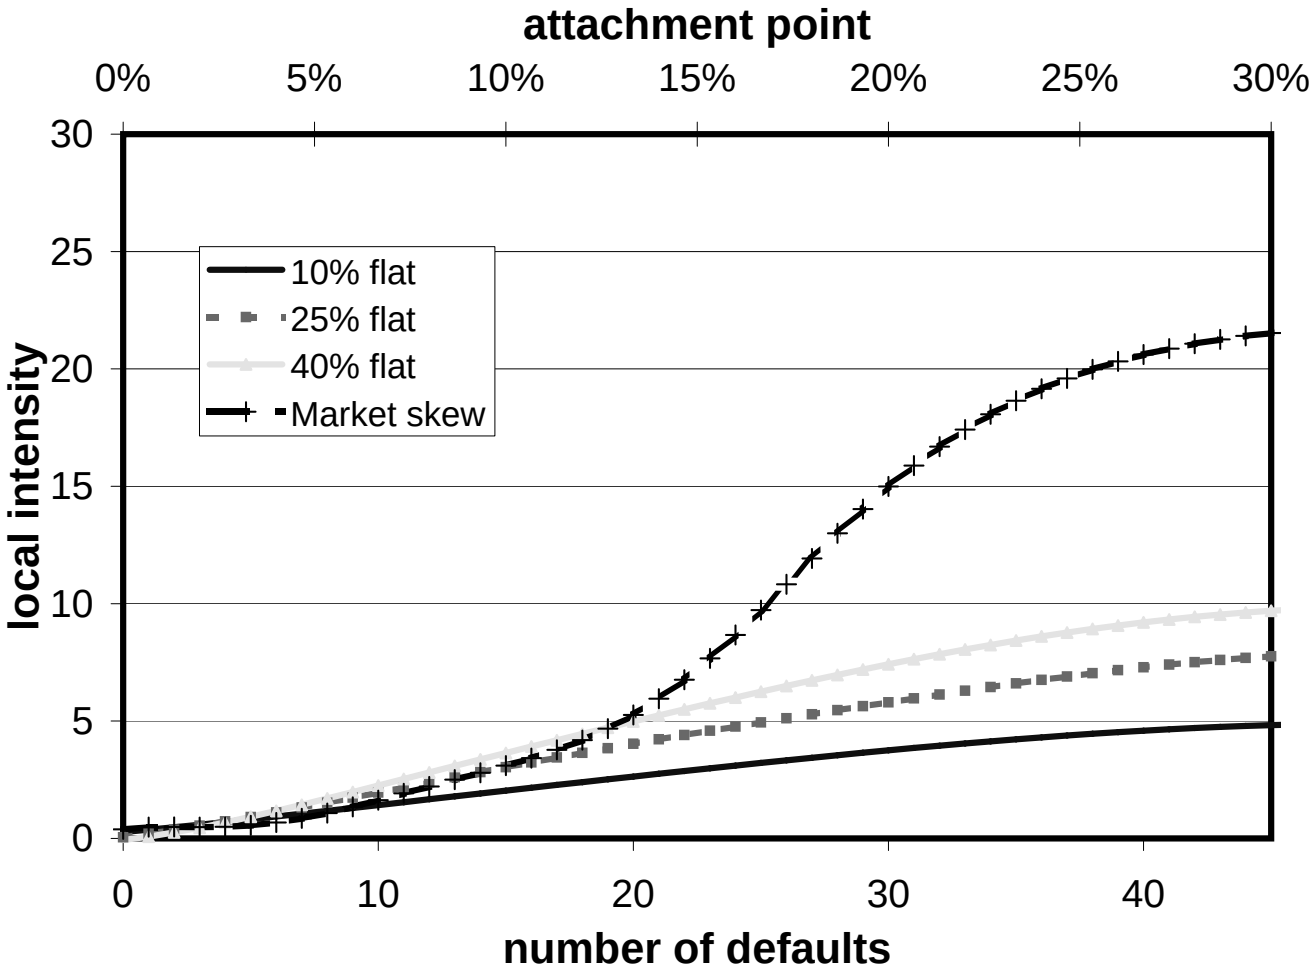

Figure 5: Dependence of the local intensity on the number of defaults at maturity for different values of flat Gaussian correlations and for the market correlations skew.

Confidential Information Copyright ©1996–2007 NumeriX LLC All rights reserved.