

ESTIMATING TRANCHE SPREADS BY LOSS PROCESS SIMULATION

Kay Giesecke
Baeho Kim

Department of Management Science and Engineering
414 Terman Center, Stanford University
Stanford, CA 94305, U.S.A.

ABSTRACT

A credit derivative is a path dependent contingent claim on the aggregate loss in a portfolio of credit sensitive securities. We estimate the value of a credit derivative by Monte Carlo simulation of the affine point process that models the loss. We consider two algorithms that exploit the direct specification of the loss process in terms of an intensity. One algorithm is based on the simulation of intensity paths. Here discretization introduces bias into the results. The other algorithm facilitates exact simulation of default times and generates an unbiased estimator of the derivative price. We implement the algorithms to value index and tranche swaps, and we calibrate the loss process to quotes on the CDX North America High Yield index.

1 INTRODUCTION

A credit derivative is a path dependent contingent claim on the aggregate loss due to default in a portfolio of credit sensitive securities such as loans, bonds or credit swaps. The loss follows a stochastic point process that jumps at a default by the magnitude of the loss. In this paper, we show how to estimate the value of a credit derivative by Monte Carlo simulation of the loss process. To illustrate this approach, we implement the simulation based valuation and market calibration for index and tranche swaps.

We adopt the top down approach described in (Giesecke and Goldberg 2005) and specify the loss process directly in terms of an intensity that represents the conditional default rate and a distribution of the loss magnitudes. We assume that the intensity follows a mean-reverting square-root jump-diffusion process whose jump term is the loss process itself. Therefore past defaults and their recoveries influence future loss dynamics: the loss process is self-affecting and incorporates the negative dependence between default and recovery rates that is observed in credit markets.

We discuss two simulation algorithms that exploit the intensity based specification of the loss process. One algo-

rithm is based on the time change theorem of (Meyer 1971), which implies that under mild conditions, a general point process can be time-changed into a standard Poisson process. The time change is given by the cumulative intensity (the compensator). We use the converse of this result and simulate default times by re-scaling Poisson arrivals. This approach requires the generation of intensity paths, which is time-consuming and introduces discretization bias. This bias can be avoided with an exact algorithm that does not require path simulation. The exact algorithm is based on the observation that given a default time T^k , the waiting time to the next default is the first arrival time of a counting process started at T^k . The conditional distribution of the waiting time can be explicitly calculated as a function of the intensity at T^k , and the waiting time can be sampled directly from this distribution.

The algorithms are valid for the entire class of affine point processes specified in (Errais, Giesecke, and Goldberg 2006). They can be adapted to simulate many other intensity based top down models, including the specifications in (Arnsdorf and Halperin 2007), (Brigo, Pallavicini, and Torresetti 2006), (Davis and Lo 2001), (Ding, Giesecke, and Tomecek 2006), (Giesecke and Tomecek 2005), (Longstaff and Rajan 2006), (Lopatin and Misirpashaev 2007) and (Tavella and Krekel 2006). In some situations simulation is preferred to the (semi-) analytical, characteristic function or PDE methods that are used in the top down modeling literature. For example, simulation facilitates the study of instruments with complex path-dependent payoff profiles such as cash collateralized debt obligations, which are difficult to analyze with other methods.

(Duffie and Singleton 1998) and (Duffie and Garleanu 2001) use similar algorithms to simulate bottom up doubly stochastic models of the individual portfolio constituents. Using a direct specification of the portfolio loss, we avoid the separate simulation of the different constituent models. This leads to a significant reduction of the computational effort in case we are interested only in derivatives on the aggregate loss but not the constituents.

This article is structured as follows. In Section 2 we recall the pricing of index and tranche swaps from (Errais, Giesecke, and Goldberg 2006). The loss process is specified in Section 3. In Section 4 we discuss the re-scaling algorithm. The exact algorithm is analyzed in Section 5. Implementation details are explained in Section 6. In Section 7 we consider the calibration to market index and tranche spreads. Section 8 concludes.

2 FAIR INDEX AND TRANCHE SPREADS

Uncertainty is modeled by a complete probability space (Ω, \mathcal{F}, P) , where P is a pricing measure with respect to a constant risk-free rate of interest $r > 0$. Consider a portfolio of n credit sensitive positions. The default process N counts the number of defaults. It has a jump of size 1 at each default time and is otherwise constant. The loss process L records the cumulative financial loss due to default. It jumps by the magnitude of the loss at each default and is otherwise constant. These processes are adapted to a right continuous and complete filtration $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ that models the evolution of investor information.

A portfolio credit derivative is a contingent claim on N or L . It allows investors to trade default insurance on the portfolio. An *index swap* is based on a portfolio whose n constituent single name swaps have common notional 1, common maturity date T and common premium payment dates (t_m) . The default leg of an index swap is a stream of payments that cover portfolio losses as they occur. Its value at time 0 is given by the discounted cumulative losses

$$D(T) = E \left[\int_0^T e^{-rt} dL_t \right]. \quad (1)$$

The premium leg of an index swap is a stream of payments at dates (t_m) . The cash flow at t_m is a fraction S of the total notional on the names that have survived until t_m . Neglecting premium accruals, the value at time 0 of the premium leg is given by

$$P(T, S) = SE \left[\sum_m e^{-rt_m} c_m (n - N_{t_m}) \right] \quad (2)$$

where c_m is the appropriate day count fraction for the period m . The fair index swap spread S is the solution $S = S(T)$ to the equation $D(T) = P(T, S)$.

Investors interested in narrower risk profiles invest in contracts based on a *tranche* of the index specified by a lower attachment point $\underline{K} \in [0, 1]$ and an upper attachment point $\bar{K} \in (\underline{K}, 1]$. The product of the difference $K = \bar{K} - \underline{K}$ and the portfolio notional n is the tranche notional.

The default leg of a tranche swap is a stream of payments that cover portfolio losses as they occur, given that the cumulative losses are larger than $\underline{K}n$ but do not exceed $\bar{K}n$.

The cumulative payments at time t , denoted U_t , are

$$U_t = (L_t - \underline{K}n)^+ - (L_t - \bar{K}n)^+.$$

The value of these payments at time 0 is

$$D(\underline{K}, \bar{K}, T) = E \left[\int_0^T e^{-rt} dU_t \right]. \quad (3)$$

This formula is analogous to formula (1) for the value of an index swap default leg. The latter can be viewed as the default leg of a tranche swap for which $\underline{K} = 0$ and $\bar{K} = 1$.

The premium leg of a tranche swap has two parts. The first part is an upfront payment, which is expressed as a fraction F of the tranche notional Kn . The second part is a stream of payments at dates (t_m) . For a tranche with $\bar{K} < 1$, the cash flow at t_m is a fraction S of the difference between the tranche notional and the tranche loss at t_m . Neglecting accruals, the value of the premium leg is given by

$$P(\underline{K}, \bar{K}, T, F, S) = FK n + SE \left[\sum_m e^{-rt_m} c_m (Kn - U_{t_m}) \right]. \quad (4)$$

For a fixed upfront payment rate F , the fair tranche spread S is the solution $S = S(\underline{K}, \bar{K}, T, F)$ to the equation $D(\underline{K}, \bar{K}, T) = P(\underline{K}, \bar{K}, T, F, S)$. Similarly, for a fixed tranche spread S , the fair tranche upfront rate F is the solution $F = F(\underline{K}, \bar{K}, T, S)$ to the equation $D(\underline{K}, \bar{K}, T) = P(\underline{K}, \bar{K}, T, F, S)$. To assess the mark-to-market or fair spread/upfront rate of an index or tranche we model the default and loss point processes.

3 DEFAULT AND LOSS PROCESSES

The default process N is a submartingale in the filtration \mathbb{F} . The Doob-Meyer theorem guarantees that N can be decomposed into the sum of a non-decreasing predictable process A starting at 0 and an \mathbb{F} -martingale. We assume that A is absolutely continuous with respect to Lebesgue measure with non-negative, \mathbb{F} -adapted *intensity* process λ . Intuitively, λ_t is the \mathcal{F}_t -conditional event arrival rate at time t . In a portfolio of n names, N is stopped at n so the intensity jumps to 0 at the n th default.

We directly specify the intensity. We assume that λ is driven by three sources of randomness: the timing of events, the loss at events and a standard Brownian motion W that evolves independently of L and is adapted to \mathbb{F} . More precisely, on the event $\{N_t < n\}$ the process λ satisfies the stochastic differential equation

$$d\lambda_t = \kappa(c - \lambda_t)dt + \sigma\sqrt{\lambda_t}dW_t + \delta dL_t \quad (5)$$

with $\lambda_0 > 0$. At an event, the loss process L jumps and so does the intensity. The jump size is the product of the loss and the sensitivity parameter $\delta \geq 0$. Each loss is drawn




Figure 1: Sample paths of the intensity (5) and the loss process generated with the re-scaling algorithm described in Section 4. A jump in the intensity represents the impact of a default. The jump size is the product of the loss at default and the sensitivity parameter $\delta = 1$. The loss at default is drawn from an independent uniform distribution on $\{0.4, 0.6, 0.8, 1\}$. The reversion rate $\kappa = 5$ and the reversion level $c = \lambda_0 = 0.7$. The volatility $\sigma = 0.2$ controls the diffusive fluctuation of the intensity between events. The portfolio contains $n = 100$ names.

independently of previous losses as well as N and W from a distribution ν on $(0, 1]$. After the event the intensity reverts back to the level c exponentially in mean at a rate $\kappa \geq 0$, with diffusive fluctuations that are driven by the Brownian motion W . The parameter $\sigma \geq 0$ describes the diffusive volatility of the intensity. If $2\kappa c \geq \sigma^2$, then $\lambda_t > 0$ almost surely. Sample paths of the loss process and the intensity are in Figure 1. The product formula shows that equation (5) has solution

$$\begin{aligned} \lambda_t = & c + (\lambda_0 - c)e^{-\kappa t} + \sigma \int_0^t e^{-\kappa(t-s)} \sqrt{\lambda_s} dW_s \\ & + \delta \int_0^t e^{-\kappa(t-s)} dL_s. \end{aligned}$$

The arrival intensity depends on the loss process L itself. It follows that L is *self-affecting*. This property captures the feedback effects of a default. Moreover, the intensity and the recovery rate are negatively correlated. Both event feedback (contagion) and negative dependence between default and recovery rates are empirically well documented phenomena, see (Jorion and Zhang 2006) and (Altman et al. 2005), respectively.

The model (5) nests several classical models. If $\sigma = 0$, then L is a compound (Hawkes 1971) process stopped at the n th arrival, which is a stopped linear birth process if $\kappa = 0$.

If $\delta = 0$, then L is a stopped compound point process that is doubly stochastic. This means that conditional on a path of W , L is a stopped inhomogeneous Poisson process. Up to the n th event, the intensity follows a (Feller 1951) diffusion, which in finance is also known as a (Cox, Ingersoll, and Ross 1985) process.

The intensity model (5) can be generalized to multiple dimensions, to include coefficients that are deterministic functions of time, and to include additional jump terms. The associated point process belongs to the family of *affine point processes*. (Errais, Giesecke, and Goldberg 2006) show that the conditional characteristic function of an affine point process is exponentially affine in the intensity, with coefficient functions satisfying a system of Riccati ordinary differential equations. A call option $E[(L_t - K)^+]$ on the loss can then be valued numerically by Fourier inversion. Call values for different strikes K and maturities t determine the fair spread on a tranche.

In this article we estimate index and tranche spreads by evaluating the expectations (1)-(4) using Monte Carlo simulation of the default and loss point processes. A replication consists of the sequence (T^k, ℓ^k) of ordered default times $0 = T^0 < T^1 < T^2 < \dots$ on $[0, T]$ and loss magnitudes ℓ^k . The associated default and loss process paths are given by the sums $N_t = \sum_{k=1}^n 1(T^k \leq t)$ and $L_t = \sum_{k=1}^n \ell^k 1(T^k \leq t)$. Below we describe alternative approaches to generate a replication. Each approach exploits the specification of the loss process through the intensity (5).

4 TIME CHANGE

Since the default process N has intensity λ , the compensator $A_t = \int_0^t \lambda_s ds$ is continuous. The intensity λ_t is strictly positive on the event $\{N_t < n\}$ and jumps to 0 at the n th default. Therefore, A is strictly increasing on $\{N_t < n\}$ and constant on $\{N_t \geq n\}$. Let $A_\infty = \lim_{t \rightarrow \infty} A_t$. For $s \in [0, A_\infty)$, define $A_s^{-1} = \inf\{t : A_t > s\}$. The process A^{-1} is continuous and strictly increasing. It follows that each A_s^{-1} is a finite \mathbb{F} -stopping time. The time change theorem of (Meyer 1971) implies that the process $N_{A^{-1}}$ is a standard Poisson process on $[0, A_\infty)$ in the right continuous, time-changed filtration \mathbb{G} generated by the stopping time sigma-fields $\mathcal{F}_{A_s^{-1}}$ for $s \in [0, A_\infty)$. Therefore, with T^k denoting the k th default time, the random variables

$$S^k = A_{T^k} = \int_0^{T^k} \lambda_s ds, \quad k = 1, 2, \dots, n \quad (6)$$

are the arrival times of a standard \mathbb{G} -Poisson process in $[0, A_\infty)$. The inter-arrival times $S^k - S^{k-1}$ are mutually independent, standard exponential random variables. Conversely (see (Giesecke and Tomecek 2005)), the k th default time is

the hitting time of A to the random level S^k : almost surely,

$$T^k = A_{S^k}^{-1} = \inf \left\{ t : \int_0^t \lambda_s ds > S^k \right\}.$$

To simulate L by re-scaling of Poisson arrivals we need to generate a trajectory of the intensity λ . If $\delta > 0$, the intensity is updated at an event so the path must be generated piece-wise between arrivals. From equation (5), the k th inter-arrival intensity h follows the Feller diffusion

$$dh_t = \kappa(c - h_t)dt + \sigma \sqrt{h_t} dW_t \quad (7)$$

started at $h_0 = \lambda_{T^{k-1}}$. The condition $2\kappa c \geq \sigma^2$ guarantees that $h_t > 0$ almost surely. In case $\sigma = 0$, h is deterministic. If $\sigma > 0$, Euler discretization of the SDE (7) can be used to approximate the path on a discrete-time grid $[0 = t_0 < t_1 < \dots < t_M = t]$ with equal spacing Δt . This scheme is easy to implement, but the probability of getting a negative value for h_{t_i} is strictly positive. An alternative is to sample directly from the transition law of h_{t_i} given $h_{t_{i-1}}$. As shown by (Cox, Ingersoll, and Ross 1985), up to a scale factor, the distribution of h_s given h_t is non-central chi-squared. For $x \geq 0$, $y, \sigma > 0$ and $0 \leq t < s$, the density $g(x; y, s - t)dx = P[h_s \in dx | h_t = y]$ is given by

$$g(x; y, s) = ae^{-(ax+by)} \left(\frac{ax}{by} \right)^{\frac{q}{2}} I_q(2\sqrt{axy}) \quad (8)$$

where I_q is the modified Bessel function of the first kind of order $q = 2\kappa c / \sigma^2 - 1$ and

$$a = 2\kappa / (\sigma^2(1 - \exp(-\kappa s))) \quad (9)$$

$$b = a \exp(-\kappa s). \quad (10)$$

The following *re-scaling algorithm* generates the k th default time and loss and the intensity at the k th default given the previous default time $T^{k-1}(\omega)$ and intensity $\lambda_{T^{k-1}}(\omega)$:

- (i) Draw a standard uniform variate $U(\omega)$.
- (ii) Generate a path $h(\omega)$ of the inter-arrival intensity with $h_0(\omega) = \lambda_{T^{k-1}}(\omega)$ and set

$$T^k(\omega) = T^{k-1}(\omega) + \inf \left\{ t : \int_0^t h_s(\omega) ds > -\log U(\omega) \right\}.$$

- (iii) Draw $\ell^k(\omega)$ from the distribution ν on $(0, 1]$.
- (iv) Set $\lambda_{T^k}(\omega) = h_{T^k - T^{k-1}}(\omega) + \ell^k(\omega)$.

Figure 1 shows a sample path of the loss process and the intensity generated with the re-scaling algorithm, and Figure 2 shows the associated path of the compensator,




Figure 2: Sample paths of the compensator A , standard Poisson arrival times S^k and default times T^k associated with the intensity and loss process paths shown in Figure 1. The compensator path indicates the response of the intensity to events.

Poisson arrivals and default times. Step (ii) is implemented as follows. The inter-arrival intensity values $h_{t_i}(\omega)$ on the discrete-time grid with spacing $\Delta t = 1/365$ are generated by sampling from the chi-squared density using the method described in Section 3.1 of (Broadie and Kaya 2006). The corresponding values of the integrated inter-arrival intensity are given by the running sums of the terms $h_{t_i}(\omega)(t_i - t_{i-1})$. The inter-arrival time $T^k(\omega) - T^{k-1}(\omega)$ is the first grid time point t_i at which the running sum exceeds $-\log U(\omega)$. If T is the horizon, we stop if $t_i > T - T^{k-1}(\omega)$. Then $T^{k-1}(\omega)$ is the last default time in $[0, T]$. The simulation environment is described in Section 6.

5 EXACT SIMULATION

Unless the diffusive volatility $\sigma = 0$, the simulation of the intensity path is computationally intensive and introduces discretization bias. It can be avoided by observing that given $\mathcal{F}_{T^{k-1}}$, the waiting time $T^k - T^{k-1}$ to the next default is the first arrival time of a doubly stochastic counting process with intensity h , started at T^{k-1} with $h_0 = \lambda_{T^{k-1}}$. The doubly stochastic property and the strong Markov property of the inter-arrival intensity h imply that the conditional survival function of the waiting time is given by

$$P[T^k - T^{k-1} > t | \mathcal{F}_{T^{k-1}}] = p(t; \lambda_{T^{k-1}}) \quad (11)$$

where for $t \geq 0$ and $y > 0$,

$$p(t; y) = E \left[e^{-\int_0^t h_s ds} \mid h_0 = y \right].$$

We observe that $p(t; y)$ is the price at time 0 of a zero coupon bond maturing at time t if the short term interest rate follows the process h and the initial short rate is $h_0 = y$. But h is governed by the Feller diffusion (7), so

$$p(t; y) = e^{\alpha(t) + \beta(t)y}. \quad (12)$$

For $\sigma = 0$, the inter-arrival intensity h is deterministic and the coefficient functions can be obtained by a simple integration. For $\sigma > 0$ the coefficient functions are given by

$$\alpha(t) = \frac{\kappa c}{\sigma^2} \left(\kappa t - 2 \log \frac{\cosh(-\frac{1}{2}\gamma t + \operatorname{atanh}(-\kappa/\gamma))}{\cosh(\operatorname{atanh}(-\kappa/\gamma))} \right)$$

$$\beta(t) = \frac{1}{\sigma^2} \left(\kappa + \gamma \tanh(-\frac{1}{2}\gamma t + \operatorname{atanh}(-\kappa/\gamma)) \right)$$

where $\gamma = \sqrt{\kappa^2 + 2\sigma^2}$; see (Cox, Ingersoll, and Ross 1985).

Given the conditional survival function $p(t; y)$ and an intensity value $\lambda_{T^{k-1}}$ we can simulate the inter-arrival time $T^k - T^{k-1}$ by the inverse method. Conditional on T^k , we must then generate the intensity λ_{T^k} in order to simulate the next inter-arrival time $T^{k+1} - T^k$. This means we need to sample from the conditional distribution of h_τ given h_0 and τ , where τ is the first arrival time of a doubly stochastic counting process with intensity h . Consider the corresponding conditional density $c(x; y, t) dx = P[h_\tau \in dx | h_0 = y, \tau = t]$. For $x \geq 0$, $y > 0$ and $t \geq 0$, by Bayes' rule we have

$$c(x; y, t) = \frac{q(x, t; y)}{-\partial_t p(t; y)} \quad (13)$$

where $q(x, t; y) dx dt = P[h_\tau \in dx, \tau \in dt | h_0 = y]$ is the joint conditional density of (h_τ, τ) , and $-\partial_t p(t; y) dt = P[\tau \in dt | h_0 = y] > 0$ is the conditional density of τ . By Bayes' rule and the doubly stochastic property of τ , we get

$$q(x, t; y) = x g(x; y, t) E \left[e^{-\int_0^t h_s ds} | h_0 = y, h_t = x \right] \quad (14)$$

where $g(x; y, t)$ is the non-central chi-squared density (8). The conditional expectation on the right hand side of equation (14) can be calculated explicitly if h follows a Feller diffusion; see (Broadie and Kaya 2006) for a formula.

The following *exact algorithm* generates the k th default time and loss and the intensity at the k th default given the previous default time $T^{k-1}(\omega)$ and intensity $\lambda_{T^{k-1}}(\omega)$:

- (i) Draw a standard uniform variate $U(\omega)$ and set

$$T^k(\omega) = T^{k-1}(\omega) + \inf \{ t : p(t; \lambda_{T^{k-1}}(\omega)) \leq U(\omega) \}.$$

- (ii) If $\sigma > 0$, draw $h_{T^k - T^{k-1}}(\omega)$ from the conditional density $c(x; \lambda_{T^{k-1}}(\omega), T^k(\omega) - T^{k-1}(\omega))$ where $c(x; y, s)$ is given by equation (13).
- (iii) Draw $\ell^k(\omega)$ from the distribution ν on $(0, 1]$.
- (iv) Set $\lambda_{T^k}(\omega) = h_{T^k - T^{k-1}}(\omega) + \ell^k(\omega)$.

Figure 3 shows a path of the loss process generated with the exact algorithm, along with the corresponding intensity values just before and at the default times.

Step (i) is implemented as follows. Given the horizon T , we first evaluate $p(T - T^{k-1}(\omega); \lambda_{T^{k-1}}(\omega))$. If this value is larger than $U(\omega)$ we stop. In this case $T^{k-1}(\omega)$ is the last default time in $[0, T]$. Otherwise, we numerically solve $p(t; \lambda_{T^{k-1}}(\omega)) - U(\omega) = 0$ for $t \in [0, T - T^{k-1}(\omega)]$ using Brent's method, which combines root bracketing, bisection and inverse quadratic interpolation; see Chapter 4 in (Brent 1973). The solution is equal to $T^k(\omega) - T^{k-1}(\omega)$.

Step (ii) is implemented by acceptance/rejection sampling from the chi-squared density $g(x; y, t)$. This method is based on the following observation. Using the results in (Broadie and Kaya 2006) we can show that for fixed $y > 0$ and $t > 0$, the conditional expectation on the right hand side of equation (14) is of the form

$$E \left[e^{-\int_0^t h_s ds} | h_0 = y, h_t = x \right] = V(x) C e^{-\xi x} \quad (15)$$

where for $d = 4\kappa c / \sigma^2$, $\gamma = \sqrt{\kappa^2 + 2\sigma^2}$ and I_ν denoting the modified Bessel function of the first kind of order ν ,

$$V(x) = \frac{I_{0.5d-1} \left(\sqrt{xy} \frac{4\gamma e^{-0.5\gamma t}}{\sigma^2(1-e^{-\gamma t})} \right)}{I_{0.5d-1} \left(\sqrt{xy} \frac{4\kappa e^{-0.5\kappa t}}{\sigma^2(1-e^{-\kappa t})} \right)}$$

and where C is a constant that is positive if $\sigma > 0$ and ξ is a positive constant given by

$$\xi = -\frac{1}{\sigma^2} \left(\frac{\kappa(1+e^{-\kappa t})}{1-e^{-\kappa t}} - \frac{\gamma(1+e^{-\gamma t})}{1-e^{-\gamma t}} \right).$$

Equations (13)-(15) imply that for fixed $y > 0$, $t > 0$ and $D = -C/\partial_t p(t; y) > 0$,

$$c(x; y, t) = DxV(x)e^{-\xi x} g(x; y, t) \leq DMg(x; y, t)$$

where $M = \sup\{xV(x)e^{-\xi x} : x \geq 0\}$. If $\sigma > 0$, we have $\gamma > \kappa \geq 0$. Hence, $V(x) \leq 1$ for $x \geq 0$ since $xe^{-0.5x}/(1-e^{-x})$ is strictly decreasing in $x \geq 0$ and $I_\nu(z)$ is strictly increasing in $z \geq 0$. Since $\xi > 0$ and $xe^{-\xi x}$ achieves its maximum at $x = 1/\xi$, we conclude that $M \in (0, \frac{1}{\xi}]$. Moreover, $V(x)$ is continuous and strictly decreasing in $x \geq 0$ if $\sigma > 0$. This implies that the maximum is always achieved but never achieved at $x > \frac{1}{\xi}$ since $xe^{-\xi x}$ is also decreasing in




Figure 3: Sample path of the loss process generated with the exact algorithm, along with the corresponding intensity values just before and at the default times. The difference between the intensity values just before and at an event is equal to the financial loss at the event. The parameter values are those used to generate Figures 1 and 2.

$x > 1/\xi$. Our acceptance/rejection implementation of Step (ii) of the exact algorithm is therefore as follows:

- Compute $M = \max\{xe^{-\xi x}V(x) : x \in (0, 1/\xi]\}$ with the initial point at $x = 1/\xi$ using the Levenberg-Marquardt algorithm.
- Draw $X(\omega)$ from the non-central chi-squared density $g(x; \lambda_{T^{k-1}}(\omega), T^k(\omega) - T^{k-1}(\omega))$.
- Draw a standard uniform variate $U(\omega)$.
- Accept $h_{T^k - T^{k-1}}(\omega) = X(\omega)$ and stop if

$$U(\omega) \leq \frac{c(X(\omega); \lambda_{T^{k-1}}(\omega), T^k(\omega) - T^{k-1}(\omega))}{DMg(X(\omega); \lambda_{T^{k-1}}(\omega), T^k(\omega) - T^{k-1}(\omega))} \\ = X(\omega)V(X(\omega))e^{-\xi X(\omega)}/M.$$

Otherwise, reject $X(\omega)$ and iterate.

The exact algorithm extends to the class of self-affecting intensity models for which the conditional inter-arrival time distribution $p(t; y)$ can be calculated and the conditional distribution of the inter-arrival intensity h is known. For example, the inter-arrival intensity h could follow a general affine jump diffusion, in which case $p(t; y)$ retains the exponentially affine form (12); see (Duffie, Pan, and Singleton 2000). Another example is a quadratic diffusion model, in which $\log p(t; y)$ is a quadratic function of y ; see (Leippold and Wu 2002). The price of this generality is that, except for special cases, the ordinary differential equations that govern the coefficients do not admit closed form solutions and the conditional distribution of h is known only through

its characteristic function. In this situation Steps (i) and (ii) are computationally intensive.

6 IMPLEMENTATION

The simulation experiments in this paper were performed on a desktop PC with an AMD Athlon 1.00 GHz processor and 960 MB of RAM, running Windows XP Professional. The codes were written in C++ and a Windows .NET application was built. The compiler used was Microsoft Visual C++ .NET version 7.1.3088. For random number generation, the evaluation of special functions and Brent's method we used the code provided in the GNU Scientific Library. Table 1 compares the performance of the two algorithms in this environment.

Table 1: Comparison of simulation algorithms. The table shows the estimated spread on a 5 year index swap with $n = 100$ names (basis points), the 95% confidence intervals estimated by bootstrapping (basis points) and the calculation time (seconds). The parameter values are $r = 0.05$, $\delta = 1$, $\kappa = 5$, $c = \lambda_0 = 0.7$, $\sigma = 0.2$, and the loss at default is drawn from an independent uniform distribution on $\{0.4, 0.6, 0.8, 1.0\}$. 100,000 replications were generated.

Algorithm	Spread	95% CI	Time
Re-scaling	58.02	[57.79, 58.25]	54.17
Exact	57.97	[57.73, 58.20]	11.11

7 CALIBRATION

We calibrate the intensity model (5) and the distribution ν of the loss at default to market index and tranche spreads/upfront rates on the CDX North America High Yield index observed on 5/11/2007. The data were obtained from Morgan Stanley. The CDX.NA.HY index has $n = 100$ names. The attachment points are 0, 10, 15, 25, 35 and 100 percent. The $[0, 10]$ and $[10, 15]$ tranches are quoted in terms of a percentage upfront fee. The other tranches are quoted in terms of a running spread. We do not use the $[35, 100]$ (super senior) tranche in the calibration.

We assume ν is uniform on $\{\ell_1, \ell_2\}$ with $0 < \ell_1 < \ell_2 < 1$. Consistent with market practice, we set the expected loss at default $\int z d\nu(z) = 0.6$. With this condition ν is specified by one parameter, ℓ_1 say. The intensity λ is specified by the parameters $\lambda_0, c, \kappa, \sigma$ and δ . The risk-free rate $r = 0.05$.

We fit the parameter vector $\theta = (\lambda_0, c, \kappa, \sigma, \delta, \ell_1)$ to the CDX.NA.HY market data. Using adaptive simulated

annealing, we numerically solve the optimization problem

$$\min_{\theta \in \Theta} \sum_i \left(\frac{\text{MarketMid}(i) - \text{Model}(i, \theta)}{\text{MarketAsk}(i) - \text{MarketBid}(i)} \right)^2$$

subject to $2\kappa c \geq \sigma^2$

where $\Theta = [0, 5]^3 \times [0, 1] \times [0, 5] \times [0.2, 0.6]$ and the sum ranges over the data points. The market mid quote MarketMid is the arithmetic average of the observed MarketAsk and MarketBid quotes. The model value Model is a function of θ ; see Sections 2 and 3.

We analyze two model specifications: a model Mod 1 which is unrestricted and a restricted model Mod 2 for which the diffusive volatility $\sigma = 0$. In Mod 2, the loss process L is a stopped compound Hawkes process with deterministic inter-arrival intensity.

We begin by calibrating to quotes for the 5Y maturity. The data and results are summarized in Table 2. Both specifications fit the data well; the unrestricted model (average absolute percentage error AAPE 1.47%) does better than the restricted model (AAPE 2.24%). The calibrated parameter values for the two models are very similar, see Table 3, and so are the implied loss distributions. Figure 4 shows the loss distribution for different horizons implied by the unrestricted model with parameters set to the calibrated values in Table 3.

Table 2: 5Y calibration data and results. The CDX.NA.HY 5Y market quotes are observed on 5/11/2007. They are in basis points except for the [0, 10] and [10, 15] tranches, which are quoted in terms of a percentage upfront fee. 100,000 replications were generated. The values in column Mod 1 refer to spreads generated by the unrestricted model while the values in column Mod 2 refer to model spreads when $\sigma = 0$. We report the minimum value of the objective function MinObj and the average absolute percentage error AAPE.

	MarketBid	MarketAsk	Mod 1	Mod 2
0-10	70.50%	70.75%	71.11%	71.48%
10-15	34.25%	34.50%	32.85%	32.74%
15-25	316.00	319.00	316.80	311.43
25-35	79.00	81.00	81.47	77.34
Index	262.85	263.10	263.46	262.97
MinObj			41.63	60.41
AAPE			1.47%	2.24%

To assess the stability of the calibrations over time we re-calibrate both specifications at different dates starting at 5/11/2007. The initial values at a given date after 5/11 are set to the optimal values obtained at the previous calibration date. Figure 5 (Figure 6) shows the calibrated parameter




Figure 4: Loss distribution for different horizons implied by the index and tranche spreads on the 5Y CDX.NA.HY index observed on 5/11/2007. We simulate the loss process from the model Mod 1 with parameters set to the calibrated values given in Table 3 and then use kernel smoothing to non-parametrically estimate the loss distribution. The mass in the tail of the distribution increases with the horizon.

Table 3: Initial and calibrated parameter values for the unrestricted model Mod 1 and the restricted model Mod 2 (5Y). We ran several calibrations with different initial values. The setting we report here had the lowest objective function value for model Mod 1.

	λ_0	c	κ	σ	δ	ℓ_1
Initial	2.50	2.50	2.50	0.50	2.50	0.40
Mod 1	0.70	1.61	2.62	0.62	2.99	0.24
Mod 2	0.75	1.60	2.58	0.00	2.94	0.24

values for the unrestricted model Mod 1 (the restricted model Mod 2) for the different dates. The calibrated parameter values for the unrestricted model are more stable than the values for the restricted model. Table 4 (Table 5) gives the calibration errors. In each calibration, the unrestricted model leads to a lower average absolute percentage error than the restricted model.

Next we examine the fit to quotes for several maturities, all observed on 5/11/2007. We get 5 additional data points per additional maturity. First we calibrate simultaneously to 5 and 7 year CDX.NA.HY index and tranche spreads. The data and results are shown in Table 6. The calibrated parameter values are given in Table 7. For both model specifications the average absolute percentage error increases somewhat, but the fits are still very reasonable. The unrestricted model (AAPE 2.35%) outperforms the re-




Figure 5: Parameter stability: calibrated parameter values for the unrestricted model Mod 1 at different calibration dates (5Y).




Figure 6: Parameter stability: calibrated parameter values for the restricted model Mod 2 at different calibration dates (5Y).

Table 4: Comparison of the calibrations of the unrestricted model Mod 1 at the different calibration dates (5Y). The calibrated parameter values are shown in Figure 5.

Date	05/11	05/14	05/15	05/16	05/17
MinObj	41.63	58.66	46.78	46.24	58.24
AAPE	1.47%	1.70%	1.26%	1.32%	1.35%

stricted model (AAPE 3.30%). For comparison of the errors, Table 8 shows the out-of-sample forecasts of the unrestricted model for 7Y index and tranche spreads/upfront rates, when calibrated to 5Y CDX.NA.HY market quotes observed on 5/11/2007. The forecasts for the 7Y quotes are very good with a total AAPE of 3.04%.

Finally we calibrate the unrestricted model Mod 1 simultaneously to 3, 5 and 7Y CDX.NA.HY index and tranche spreads. The data and results are shown in Table 9. The calibrated parameter values are given in Table 10. We performed the calibration with and without a constraint on the expected loss at default (in the calibrations discussed above we set the expected loss to 0.6). With the constraint, we

Table 5: Comparison of the calibrations of the restricted model Mod 2 at the different calibration dates (5Y). The calibrated parameter values are shown in Figure 6.

Date	05/11	05/14	05/15	05/16	05/17
MinObj	60.41	73.43	54.29	43.88	65.23
AAPE	2.24%	2.04%	1.57%	1.45%	1.92%

got an AAPE of 6.83%. When the constraint is relaxed, we obtain an additional parameter ℓ_2 to fit and the AAPE improves to 5.93%.

8 CONCLUSION

We describe two algorithms to simulate a loss point process that is specified in terms of an intensity. The re-scaling algorithm requires the simulation of intensity paths. Path simulation is time-consuming and introduces bias into the results. The exact algorithm avoids path simulation and generates an unbiased estimate of the derivative price. While the exact algorithm requires much less time per replication than the re-scaling algorithm, speed may still be an issue for some applications. An example is model calibration, which requires a large number of price evaluations at different parameters values. The bottleneck in the exact algorithm is the numerical evaluation of the inverse of the conditional distribution of the inter-arrival time. At present we are investigating a thinning algorithm that avoids this step.

ACKNOWLEDGMENTS

This research was supported by a grant from JP Morgan Chase’s Academic Outreach Program, for which we are very grateful. We would like to thank Lisa Goldberg and Jeremy Staum for useful comments and suggestions. We are grateful to Morgan Stanley for providing the index and tranche market data.

Table 6: 5+7Y calibration data and results. The CDX.NA.HY 5Y and 7Y market quotes are observed on 5/11/2007. They are in basis points except for the 0-10 and 10-15 percent tranches, which are quoted in terms of a percentage upfront fee. The first 5 rows are for the 5Y index while rows 6-10 are for the 7 Y index. 100,000 replications were generated. The values in column Mod 1 refer to spreads generated by the unrestricted model while the values in column Mod 2 refer to model spreads when $\sigma = 0$. We report the minimum value of the objective function MinObj and the average absolute percentage error AAPE.

	MarketBid	MarketAsk	Mod 1	Mod 2
0-10	70.50%	70.75%	71.69%	72.00%
10-15	34.25%	34.50%	33.44%	33.47%
15-25	316.00	319.00	316.13	309.93
25-35	79.00	81.00	78.45	72.79
Index	262.85	263.10	263.94	262.75
0-10	80.13%	80.38%	81.23%	81.49%
10-15	55.50%	55.75%	55.17%	55.40%
15-25	582.00	587.00	580.25	584.63
25-35	180.00	183.00	206.41	207.98
Index	307.50	307.75	307.53	308.91
MinObj			136.12	192.82
AAPE			2.35%	3.30%

REFERENCES

Altman, E., B. Brady, A. Resti, and A. Sironi. 2005. The link between default and recovery rates: Theory, empirical evidence and implications. *Journal of Business* 78 (6): 2203–2227.

Arnsdorf, M., and I. Halperin. 2007. BSLP: Markovian bivariate spread-loss model for portfolio credit derivatives. Working Paper, Quantitative Research J.P. Morgan.

Brent, R. 1973. *Algorithms for minimization without derivatives*. Englewood Cliffs: Prentice Hall.

Table 7: 5+7Y calibration: initial and calibrated parameter values for the unrestricted model Mod 1 and the restricted model Mod 2. We used the same initial values as in the 5Y calibration.

	λ_0	c	κ	σ	δ	ℓ_1
Initial	2.50	2.50	2.50	0.50	2.50	0.40
Mod 1	0.13	1.77	2.10	0.69	2.44	0.25
Mod 2	0.12	1.76	1.96	0.00	2.32	0.24

Table 8: Out-of-sample forecasts of the unrestricted model Mod 1 for 7Y index and tranche spreads/upfront rates. Mod 1 is calibrated to 5Y CDX.NA.HY market quotes observed on 5/11/2007 (first 5 rows; calibrated model values are in the last column). The model forecast values for the 7Y quotes are given in rows 6-10, last column. 100,000 replications were generated. We report the total average absolute percentage error AAPE for the 5 and 7Y data (calibration and forecast).

		MarketBid	MarketAsk	Mod 1
5Y	0-10	70.50%	70.75%	71.11%
	10-15	34.25%	34.50%	32.85%
	15-25	316.00	319.00	316.80
	25-35	79.00	81.00	81.47
	Index	262.85	263.10	263.46
7Y	0-10	80.13%	80.38%	81.79%
	10-15	55.50%	55.75%	52.49%
	15-25	582.00	587.00	531.61
	25-35	180.00	183.00	186.26
	Index	307.50	307.75	296.00
AAPE				3.04%

Brigo, D., A. Pallavicini, and R. Torresetti. 2006. Calibration of CDO tranches with the dynamical generalized-Poisson loss model. Working Paper, Banca IMI.

Broadie, M., and O. Kaya. 2006. Exact simulation of stochastic volatility and other affine jump diffusion processes. *Operations Research* 54 (2): 217–231.

Cox, J., J. Ingersoll, and S. Ross. 1985. A theory of the term structure of interest rates. *Econometrica* 53:385–408.

Davis, M., and V. Lo. 2001. Modeling default correlation in bond portfolios. In *Mastering Risk Volume 2: Applications*, ed. C. Alexander, 141–151. Prentice Hall.

Ding, X., K. Giesecke, and P. Tomecek. 2006. Time-changed birth processes and multi-name credit. Working Paper, Stanford University.

Duffie, D., and N. Garleanu. 2001. Risk and valuation of collateralized debt obligations. *Financial Analysts Journal* 57 (1): 41–59.

Duffie, D., J. Pan, and K. Singleton. 2000. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68:1343–1376.

Duffie, D., and K. J. Singleton. 1998. Simulating correlated defaults. Working Paper, GSB, Stanford University.

Errais, E., K. Giesecke, and L. Goldberg. 2006. Pricing credit from the top down with affine point processes. Working Paper, Stanford University.

Feller, W. 1951. Two singular diffusion problems. *Annals of Mathematics* 54:173–182.

Giesecke, K., and L. Goldberg. 2005. A top down approach to multi-name credit. Working Paper, Stanford University.

Table 9: 3+5+7Y calibration data and results for the unrestricted model Mod 1 without a constraint on the expected loss at default. The CDX.NA.HY 3Y, 5Y and 7Y market quotes are observed on 5/11/2007. They are in basis points except for the 0-10 and 10-15 percent tranches, which are quoted in terms of a percentage upfront fee. 100,000 replications were generated. We report the minimum value of the objective function MinObj and the average absolute percentage error AAPE.

		MarketBid	MarketAsk	Mod 1
3Y	0-10	45.63%	45.88%	45.03%
	10-15	7.75%	8.00%	7.51%
	15-25	51.00	55.00	54.94
	25-35	12.00	14.00	7.77
	Index	184.70	184.95	189.21
5Y	0-10	70.50%	70.75%	70.25%
	10-15	34.25%	34.50%	32.10%
	15-25	316.00	319.00	310.74
	25-35	79.00	81.00	79.26
	Index	262.85	263.10	258.69
7Y	0-10	80.13%	80.38%	80.25%
	10-15	55.50%	55.75%	53.56%
	15-25	582.00	587.00	580.92
	25-35	180.00	183.00	218.18
	Index	307.50	307.75	308.04
	MinObj			930.19
	AAPE			5.93%

Giesecke, K., and P. Tomecek. 2005. Dependent events and changes of time. Working Paper, Stanford University.

Hawkes, A. G. 1971. Spectra of some self-exciting and mutually exciting point processes. *Biometrika* 58 (1): 83–90.

Jorion, P., and G. Zhang. 2006. Good and bad credit contagion: Evidence from credit default swaps. *Journal of Financial Economics*, Forthcoming.

Leippold, M., and L. Wu. 2002. Asset pricing under the quadratic class. *Journal of Financial and Quantitative Analysis* 37:271–295.

Table 10: 3+5+7Y calibration: initial and calibrated parameter values for the unrestricted model Mod 1 without a constraint on the expected loss at default.

	λ_0	c	κ	σ	δ	ℓ_1	ℓ_2
Initial	2.50	2.50	2.50	0.50	2.50	0.40	0.80
Mod 1	0.03	1.56	2.07	0.51	2.66	0.34	0.84

Longstaff, F., and A. Rajan. 2006. An empirical analysis of collateralized debt obligations. Forthcoming, *Journal of Finance*.

Lopatin, A., and T. Misirpashaev. 2007. Two-dimensional Markovian model for dynamics of aggregate credit loss. Working Paper, Numerix.

Meyer, P.-A. 1971. Démonstration simplifiée d’un théorème de Knight. In *Séminaire de Probabilités V, Lecture Note in Mathematics 191*, 191–195. Springer-Verlag Berlin.

Tavella, D., and M. Krekel. 2006. Pricing nth-to-default credit derivatives in the PDE framework. Working Paper, Octanti Associates.

AUTHOR BIOGRAPHIES

KAY GIESECKE is an assistant professor in the Department of Management Science and Engineering at Stanford University. He has previously held a position in the School of Operations Research and Information Engineering at Cornell University. His research is in financial engineering, particularly credit risk modeling. His email is giesecke@stanford.edu and his web page is <www.stanford.edu/~giesecke>.

BAEHO KIM is a Ph.D. candidate in the Department of Management Science and Engineering at Stanford University. His research interests include stochastic modeling and computational techniques in financial engineering. His email is baehokim@stanford.edu and web page is <www.stanford.edu/~baehokim>.