OPERATIONS RESEARCH
Vol. 59, No. 1, January–February 2011, pp. 32–49
ISSN 0030-364X | EISSN 1526-5463 | 11 | 5901 | 0032

DOI 10.1287/opre.1100.0864
© 2011 INFORMS

Risk Analysis of Collateralized Debt Obligations

Kay Giesecke
Department of Management Science and Engineering, Stanford University,
Stanford, California 94305, giesecke@stanford.edu

Baeho Kim
Korea University Business School, Anam-dong, Seongbuk-gu,
Seoul 136-701, Korea, baehokim@korea.ac.kr

Collateralized debt obligations, which are securities with payoffs that are tied to the cash flows in a portfolio of defaultable assets such as corporate bonds, play a significant role in the financial crisis that has spread throughout the world. Insufficient capital provisioning due to flawed and overly optimistic risk assessments is at the center of the problem. This paper develops stochastic methods to measure the risk of positions in collateralized debt obligations and related instruments tied to an underlying portfolio of defaultable assets. It proposes an adaptive point process model of portfolio default timing, a maximum likelihood method for estimating point process models that is based on an acceptance/rejection resampling scheme, and statistical tests for model validation. To illustrate these tools, they are used to estimate the distribution of the profit or loss generated by positions in multiple tranches of a collateralized debt obligation that references the CDX High Yield portfolio and the risk capital required to support these positions.

Subject classifications: correlated default risk; collateralized debt obligation; portfolio credit derivative; actual measure; point process; intensity; resampling; thinning; acceptance/rejection sampling; exact simulation.
Area of review: Financial Engineering.
History: Received June 2009; revisions received January 2010, March 2010; accepted March 2010. Published online in *Articles in Advance* December 30, 2010.

1. Introduction

The financial crisis highlights the need for a holistic, objective, and transparent approach to accurately measuring the risk of investment positions in portfolio credit derivatives such as collateralized debt obligations (CDOs). Portfolio credit derivatives are securities whose payoffs are tied, often through complex schemes, to the cash flows in a portfolio of credit instruments such as corporate bonds, loans, or mortgages. They facilitate the trading of insurance against the default losses in the portfolio. An investor providing the insurance is exposed to the default risk in the portfolio.

There is an extensive literature devoted to the valuation and hedging of portfolio derivatives.[1] The basic valuation problem is to estimate the price of default insurance, i.e., the arbitrage-free value of the portfolio derivative at contract inception. This value is given by the expected discounted derivative cash flows relative to a risk-neutral pricing measure. After inception, the derivative position must be marked to market; that is, the value of the derivative under current market conditions must be determined. The basic hedging problem is to estimate the sensitivities of the derivative value to changes of the default risk of the portfolio constituents. These sensitivities determine the amount of constituent default insurance to be bought or sold to neutralize the derivative price fluctuations due to small changes of the constituent risks.

The valuation and hedging problems are distinct from the risk analysis problem, which is to measure the exposure of the derivative investor, who provides default insurance, to potential payments due to defaults in the portfolio. More precisely, the goal is to estimate the distribution of the investor's cumulative cash flows over the life of the contract. The distribution is taken under the actual measure describing the empirical likelihood of events, rather than a risk-neutral pricing measure. The distribution describes the risk/reward profile of a portfolio derivative position and is the key to risk management applications. For example, it allows the investor or regulator to determine the amount of risk capital required to support a position. The financial crisis indicates the significance of these applications and the problems associated with the traditional rating based analysis of portfolio credit derivative positions; see SEC (2008).

The risk analysis problem has been largely ignored in the academic literature. This paper provides stochastic methods to address this problem. It makes several contributions. First, it develops a maximum likelihood approach to estimating point process models of portfolio default timing from historical default experience. Second, it devises statistical tests to validate a fitted model. Third, it formulates, fits and tests an adaptive point process model of portfolio default timing, and it demonstrates the utility of the estimation and validation methods on this model. Fourth, it addresses important risk management applications, including the estimation of profit and loss

distributions for positions in multiple tranches of a CDO. These distributions quantify and differentiate the risk exposure of alternative investment positions and the impact of complex contract features. They are preferable to agency ratings, which are often based on the first moment only.

Estimating a stochastic point process model of portfolio default timing under the actual probability measure presents unique challenges. Most importantly, inference must be based on historical default timing data rather than market derivative pricing data. However, the default history of the reference portfolio underlying the credit derivative is often unavailable, so direct inference is usually not feasible. We confront this difficulty by developing an acceptance/rejection resampling scheme that allows us to generate alternative portfolio default histories from the available economy-wide default timing data. These histories are then used to construct maximum likelihood estimators for a portfolio point process model that is specified in terms of an intensity process. A time-scaling argument leads to testable hypotheses for the fit of a model.

The resampling approach is predicated on a top-down formulation of the portfolio point process model. This formulation has become popular in the credit derivatives pricing literature. Here, the point process intensity is specified without reference to the portfolio constituents. In our risk analysis setting, the combination of top-down formulation and resampling based inference leads to low-dimensional estimation, validation, and prediction problems that are highly tractable and fast to address even for the large portfolios that are common in practice. Alternative bottom-up formulations in Das et al. (2007), Delloye et al. (2006), and Duffie et al. (2009) require the specification and estimation of default timing models for the individual portfolio constituent securities. They allow one to incorporate firm-specific data into the estimation, but they lead to high-dimensional computational problems.

To demonstrate the effectiveness of the resampling approach and the appropriateness of the top-down formulation, we develop and fit an adaptive intensity model. This model extends the classical (Hawkes 1971) model by including a state-dependent drift coefficient in the intensity dynamics. The state-dependent drift involves a reversion level and speed that are proportional to the intensity at the previous event. While this specification is as tractable as the Hawkes model, it avoids the constraints imposed by the constant Hawkes reversion level and speed. This helps to better fit the regime-dependent behavior of empirical default rates. In- and out-of-sample tests show that our adaptive model indeed captures the clustering in the default arrival data. Because it can capture the default clustering better than other models that have been used in practice, our model leads to more realistic estimates of the probabilities of losses to senior tranches.

The rest of this paper is organized as follows. Section 2 develops the resampling approach to point process model estimation and validation. Section 3 formulates and analyzes the adaptive point process model and uses the resampling approach to fit it. Section 4 applies the fitted model to the risk analysis of synthetic CDOs, while §5 analyzes the risk of cash CDOs. Section 6 concludes. An electronic companion to this paper is available as part of the online version that can be found at http://or.journal.informs.org/.

2. Resampling Based Inference

This section develops a resampling approach to estimating a stochastic point process model of default timing. It also provides a method to validate the estimators.

2.1. Preliminaries and Problem

The uncertainty in the economy is modeled by a complete probability space $(\Omega, \mathcal{F}, P)$, where P is the actual (statistical) probability measure. The information flow of investors is described by a right-continuous and complete filtration $\mathbb{F} = (\mathcal{F}_t)_{t\geq 0}$.

Consider a nonexplosive counting process N with event stopping times $0 < T_1 < T_2 < \cdots$. The T_n represent the ordered default times in a *reference portfolio* of firms, with the convention that a defaulted name is replaced with a name that has the same characteristics as the defaulter. A portfolio credit derivative is a security with cash flows that depend on the financial loss due to default in the reference portfolio.

Suppose N has a strictly positive intensity λ such that $\int_0^t \lambda_s\, ds < \infty$ almost surely. The intensity represents the conditional mean default rate in the sense that $E(N_{t+\Delta} - N_t \mid \mathcal{F}_t) \approx \lambda_t \Delta$ for small $\Delta > 0$. This means that $N - \int_0^{\cdot} \lambda_s\, ds$ is a local martingale relative to P and $\mathbb{F}$. The process followed by λ determines the distribution of N. It is the modeling primitive and is specified without reference to the constituent firms.

Our goal is to estimate a given model of λ. If the data consist of a path of N, then this is a classical statistical problem; see Ogata (1978), for example. However, a path of N is rarely available, so direct inference is typically not feasible. Instead, the data consist of a path of the *economy-wide* default process N^* generated by the default stopping times $0 < T_1^* < T_2^* < \cdots$ in the universe of names. We develop a resampling approach to estimating λ from the realization of N^*. The basic idea is to generate alternative portfolio default histories from N^* and to estimate λ from these histories.

2.2. Acceptance/Rejection Resampling

We propose to generate paths of N from the realization of N^* by acceptance/rejection sampling. Here, we randomly select an event time of N^* as an event time of N with a certain conditional probability. The following basic observation, whose proof is found in online Appendix A, provides the foundation of this mechanism and the justification of our estimation approach. It also facilitates the design of tests to evaluate the approach.

PROPOSITION 2.1. *Let Z^* be a predictable process with values in $[0,1]$. Select an economy-wide event time T_n^* with probability $Z_{T_n^*}^*$. If the economy-wide default process N^* has intensity λ^*, then the counting process of the selected times has intensity $Z^*\lambda^*$.*

Proposition 2.1 states that the counting process obtained by thinning N^* according to Z^* has intensity given by the product of the *thinning process* Z^* and the intensity of N^*. The specification of Z^* is subject to a mild predictability condition: the selection probability at T_n^* can depend only on information accumulated up to but not including time T_n^*. Proposition 2.1 is related to the construction of a marked point process from a Poisson random measure through a state-dependent thinning mechanism in Proposition 3.1 of Glasserman and Merener (2003). It is a generalization of the classical, state-independent thinning scheme for the Monte Carlo simulation of time-inhomogeneous Poisson processes proposed by Lewis and Shedler (1979).

We generate paths of the portfolio default process N by thinning N^* with a fixed process Z^*; see Algorithm 1. Each of these paths represents an alternative event sequence, or portfolio default history, assuming a defaulter is replaced with a name that has the same characteristics as the defaulter. Proposition 2.1 implies that the events in each alternative sequence arrive with intensity $\lambda = Z^*\lambda^*$. This justifies the estimation of λ from the alternative paths of the portfolio default process N.

Algorithm 1 (Acceptance/rejection resampling)
Generating a sample path of the portfolio default process N from the realization of the economy-wide default process N^* over the sample period $[0,\tau]$.

1: Initialize $m \leftarrow 0$.
2: **for** $n = 1$ to N_τ^* **do**
3: Draw $u \sim U(0,1)$.
4: **if** $u \leqslant Z_{T_n^*}^*$ **then**
5: Assign $T_{m+1} \leftarrow T_n^*$ and update $m \leftarrow m+1$.
6: **end if**
7: **end for**

It is important to note that this estimation approach does not require the specification of a model for the economy-wide intensity λ^*. We need only to formulate a model for the thinning process Z^*, to which we turn next.

2.3. Thinning Process Specification

The random measure argument behind Proposition 3.1 in Giesecke et al. (2011) can be adapted to show that Z^* takes the form

$$Z_t^*(\omega) = \lim_{\epsilon \to 0} \frac{E(N_{t+\epsilon} - N_t \mid \mathcal{F}_t)}{E(N_{t+\epsilon}^* - N_t^* \mid \mathcal{F}_t)}(\omega) \tag{1}$$

in all those points $(\omega, t) \in \Omega \times (0,\infty]$ where the limit exists.[2] The quotient on the right side of Equation (1), which is taken to be zero when the denominator vanishes, represents the conditional probability at time t that the next defaulter is a reference name, given that a default occurs in the economy by time $t+\epsilon$.

We specify Z^* nonparametrically, guided by formula (1). Intuitively, Z^* must reflect the relation between the issuer composition of the economy and the issuer composition of the reference portfolio. We propose to describe the issuer composition in terms of the credit ratings of the respective constituent issuers. In this case, N^* is identified with the default process in the universe of rated issuers.

The use of ratings does not limit the applicability of our specification, because the reference portfolios of most derivatives consist of rated names only. Moreover, the rating agencies maintain extensive and accessible databases that record credit events including defaults in the universe of rated names, and further attributes associated with these events, such as recovery rates. The agencies have maintained a high firm coverage ratio throughout the sectors of the economy, and therefore the universe of rated names is a reasonable representation of the firms in the economy.

Let $[0,\tau]$ be the sample period, with τ denoting the (current) analysis time. Let R be the set of rating categories, $X_\tau(\rho)$ be the number at time τ of reference firms with rating $\rho \in R$, and $X_t^*(\rho)$ be the number at time $t \in [0,\tau]$ of ρ-rated firms in the universe of rated names. The number of firms $X^*(\rho)$ is an adapted process. It varies through time because issuers enter and exit the universe of rated names. Exits can be due to mergers or privatizations, for example. We assume that the thinning process Z^* is equal to the predictable projection of the process defined for times $0 < t \leqslant \tau$ by[3]

$$\frac{X_\tau(\rho^*_{N_{t-}^*+1})}{X_{t-}^*(\rho^*_{N_{t-}^*+1})} 1_{\{X_{t-}^*(\rho^*_{N_{t-}^*+1})>0\}}, \tag{2}$$

where $\rho_n^* \in \mathcal{F}_{T_n^*}$ is the rating at the time of default of the nth defaulter in the economy.[4] We require that $X_\tau(\rho) \leqslant X_t^*(\rho)$ for all $t \leqslant \tau$ and $\rho \in R$. Because $X_\tau(\rho)$ is a fixed integer prescribed by the portfolio composition and $X_{t-}^*(\rho)$ is predictable for fixed ρ,

$$Z_t^* = \sum_{\rho \in R_{t-}^*} \frac{X_\tau(\rho)}{X_{t-}^*(\rho)} P(\rho^*_{N_{t-}^*+1} = \rho \mid \mathcal{F}_{t-}) \tag{3}$$

almost surely, where R_t^* is the set of rating categories $\rho \in R$ for which $X_t^*(\rho) > 0$. Formula (3) suggests to interpret the value Z_t^* as the conditional "empirical" probability that the next defaulter is a reference name. This conditional probability respects the ratings of the reference names, as $P(\rho^*_{N_{t-}^*+1} = \rho \mid \mathcal{F}_{t-})$ is the conditional probability that the next defaulter has rating ρ. Our estimator $\nu_t^*(\rho)$ of this latter conditional probability is based on the ratings of the defaulters in $[0,t)$. For $\rho \in R$, it is given by

$$\nu_t^*(\rho) = \frac{\sum_{n=1}^{N_{t-}^*} 1_{\{\rho_n^*=\rho\}} + \alpha}{\sum_{n=1}^{N_{t-}^*} 1_{\{\rho_n^* \in R_{t-}^*\}} + \alpha |R_{t-}^*|} 1_{\{\rho \in R_{t-}^*\}}, \tag{4}$$

where $\alpha \in (0, 1]$ is an additive smoothing parameter guaranteeing that $\nu_t^*(\rho)$ is well defined for $t < T_1^*$. For $\alpha = 0$, Equation (4) defines the empirical rating distribution, which treats the observations $\rho_1^*, \ldots, \rho_{N_{t-}^*}^*$ as independent samples from a common distribution and ignores all other information contained in $\mathcal{F}_{t-}$. Our implementation assumes $\alpha = 0.5$, a value that can be justified on Bayesian grounds; see Box and Tiao (1992, pp. 34–36).[5]

2.4. Likelihood Estimators and Fitness Tests

Based on a collection of resampling paths $\{N(\omega_i)\colon i \leqslant I\}$ of the portfolio default process N generated by Algorithm 1, we estimate a model $\lambda = \lambda^\theta$, where $\theta \in \Theta$ is a parameter vector and Θ is the set of admissible parameters. The paths of N induce intensity paths $\{\lambda^\theta(\omega_i)\colon i \leqslant I, \theta \in \Theta\}$. We fit θ by solving the log-likelihood problem[6]

$$\sup_{\theta\in\Theta} \int_0^\tau \sum_{i=1}^{I} (\log \lambda_{s-}^\theta(\omega_i)\, dN_s(\omega_i) - \lambda_s^\theta(\omega_i)\, ds). \tag{5}$$

The adequacy of the fitted intensity as a model of portfolio defaults depends on the effectiveness of the resampling procedure and the appropriateness of the parametric intensity specification. More precisely, it depends on how well the alternative resampling scenarios generated by Z^* capture the actual default clustering in the reference portfolio, and how well the fitted model for λ replicates these clusters. We require statistical tests to assess this. The following result allows us to design such tests.

Proposition 2.2. *Suppose that $Z_t^* > 0$, almost surely. Then the economy-wide default process N^* is a standard Poisson process under a change of time defined by*

$$A_t^* = \int_0^t \frac{1}{Z_s^*}\, \lambda_s\, ds. \tag{6}$$

Proposition 2.2, whose proof is found in online Appendix A in the electronic companion, expresses the compensator A^* to N^* in terms of Z^* and λ and states that this compensator can be used to time-scale N^* into a standard Poisson process. We evaluate the joint specification of Z^* and λ by testing whether the fitted (mean) paths of these processes generate a realization of A^* that time-scales the observed T_n^* into a standard Poisson sequence. The Poisson property can be tested with a battery of tests.

2.5. Portfolios Without Replacement

The intensity model λ estimated from the resampling scenarios generated by Algorithm 1 is based on a portfolio with replacement of defaulters. This is without loss of generality because we can extend the reach of the fitted model to portfolios without replacement.

Consider event times T_n of N over some interval $[\tau, H]$, where $H > \tau$. The event times T_n' of the portfolio default process without replacement, N', can be obtained from the T_n by removing event times due to replaced defaulters. This is done by thinning. To formalize this, let $X_t(\rho)$ be the number of ρ-rated reference names at time $t \geqslant \tau$, assuming defaulters are not replaced. Thus, for fixed ρ, $X_t(\rho) \leqslant X_\tau(\rho)$ almost surely for every $t \geqslant \tau$. For fixed ρ, the process $X(\rho)$ decreases and vanishes when all ρ-rated reference names are in default. It suffices to specify Z, the thinning process for N, at the event times $T_n \geqslant \tau$ of N. Motivated by formula (3), we suppose that

$$Z_{T_n} = \sum_{\rho\in R_\tau} \frac{X_{T_n-}(\rho)}{X_\tau(\rho)} P(\rho_n = \rho \mid \mathcal{F}_{T_n-}), \tag{7}$$

where R_τ is the set of rating categories $\rho \in R$ for which $X_\tau(\rho) > 0$, and ρ_n is the rating of the firm defaulting at T_n. Note that the thinning probability (7) vanishes when all firms in the portfolio are in default. We estimate the conditional distribution $P(\rho_n = \cdot \mid \mathcal{F}_{T_n-})$ by the smoothed empirical distribution ν of the ratings of the defaulters in the resampling scenarios $\{N(\omega_i)\colon i \leqslant I\}$, where

$$\nu(\rho) = \frac{\sum_{i=1}^{I} \sum_{n=1}^{N_\tau(\omega_i)} 1_{\{\rho_n(\omega_i)=\rho\}} + \alpha}{\sum_{i=1}^{I} \sum_{n=1}^{N_\tau(\omega_i)} 1_{\{\rho_n(\omega_i)\in R_\tau\}} + \alpha|R_\tau|} 1_{\{\rho\in R_\tau\}} \tag{8}$$

for an additive smoothing parameter $\alpha \in [0, 1]$; see formula (4). This estimator treats the observations $\rho_n(\omega_i)$ of all paths ω_i as independent samples from a common distribution. Algorithm 2 summarizes the steps required to generate N' from N.

Algorithm 2 (Replacement thinning)
Generating a sample path of the portfolio default process N' without replacement from a path of the portfolio default process N with replacement over $[\tau, H]$, for a horizon $H > \tau$.

Initialize $m \leftarrow 0$ as we define $T_0' = \tau$, and
    set $Y(\rho) \leftarrow X_\tau(\rho)$ for all $\rho \in R$.
for $n = N_\tau + 1$ to N_H
  Draw $u \sim U(0, 1)$.
  if $u \leqslant Z_{T_n}$ **then**
    Assign $T_m' \leftarrow T_n$ and update $m \leftarrow m + 1$.
    Draw $\rho_m \sim \nu'$, where

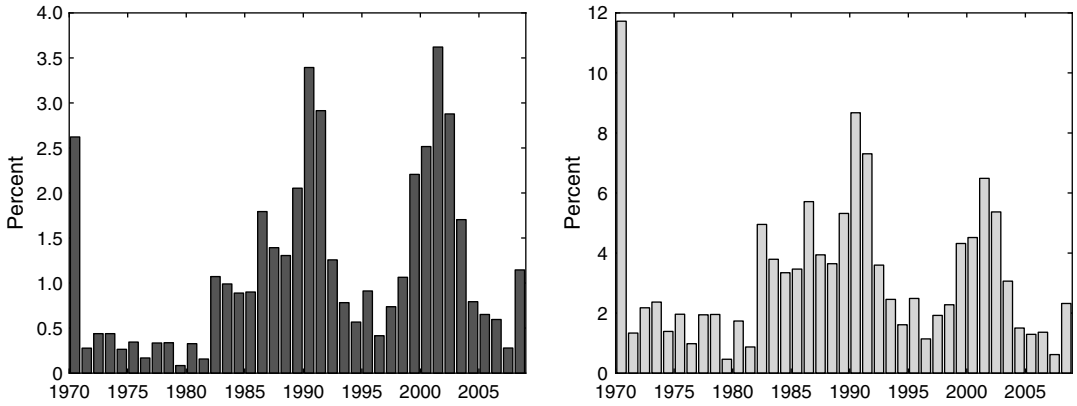
$$\nu'(\rho) = \frac{Y(\rho)}{X_\tau(\rho)} \nu(\rho)/Z_{T_n}, \quad \rho \in R. \tag{9}$$

    Update $Y(\rho_m) \leftarrow Y(\rho_m) - 1$.
  end if
end for

3. An Adaptive Intensity Model

This section demonstrates the effectiveness of the resampling approach. It formulates, fits, and evaluates an adaptive intensity model for the CDX High Yield portfolio, which is a standard portfolio that is referenced by a range of credit derivatives. The fitted intensity model will be used in §§4 and 5 to analyze the risk of such derivatives.

Figure 1. Left panel: Annual default rate, relative to the number of rated names at the beginning of a year, in the universe of Moody's rated corporate issuers in any year between 1970 and 2008, as of 11/7/2008. Right panel: Mean annual default rate for the CDX.HY6 portfolio for $I = 10$ K resampling scenarios.



Source. Moody's Default Risk Service.

3.1. Resampling Scenarios

We begin by examining the resampling scenarios of the default process of the CDX High Yield Series 6 portfolio, or CDX.HY6. We adopt the rating system of Moody's, a leading rating agency.[7] The reference portfolio consists of 2 Baa, 45 Ba, 39 B and 14 C rated firms. The realization of N^* comes from Moody's Default Risk Service and covers all Moody's rated issuers between 1/1/1970 and 11/7/2008. We observe a total of 1447 defaults.

Figure 1 shows the annual economy-wide default rate between 1970 and 2008, along with the mean default rate for the CDX.HY6 portfolio, obtained from Algorithm 1.[8] Portfolio defaults cluster heavily, indicating the positive dependence between the defaults of reference names. The excessive default rate in 1970 is due to an exceptional cluster of 24 railway defaults on June 21, 1970. Other major event clusters are related to the 1987 crash and the burst of the Internet bubble around 2001. These clusters translate into substantial fluctuations of the empirical portfolio default rate that the model portfolio intensity λ must replicate. This calls for special model features.

3.2. Intensity Specification

Our specification of λ is informed by the results of the empirical analysis in Azizpour et al. (2008b), which is based on roughly the same historical default data used here. Using in- and out-of-sample tests, Azizpour et al. (2008b) found that prediction of economy-wide default activity based on past default timing tends to be very effective. This motivates the formulation of a parsimonious portfolio intensity model whose conditioning information set is given by the past default history.

We assume that λ_t is a function of the path of N over $[0, t]$. More specifically, we propose that λ evolves through time according to the equation

$$d\lambda_t = \kappa_t(c_t - \lambda_t)\,dt + dJ_t, \tag{10}$$

where $\lambda_0 > 0$ is the initial intensity value, $\kappa_t = \kappa\lambda_{T_{N_t}}$ is the decay rate, $c_t = c\lambda_{T_{N_t}}$ is the reversion level, and J is a response jump process given by

$$J_t = \sum_{n\geqslant 1} \max(\gamma, \delta\lambda_{T_n^-}) 1_{\{T_n \leqslant t\}}. \tag{11}$$

The quantities $\kappa > 0$, $c \in (0, 1)$, $\delta > 0$ and $\gamma \geqslant 0$ are parameters. We denote by θ the vector $(\kappa, c, \delta, \gamma, \lambda_0)$. We give a sufficient condition guaranteeing that $N_t < \infty$ almost surely, for all t. This condition relates the reversion level parameter c to the parameter δ, which controls the magnitude of a jump of the intensity at an event. The proof can be found in online Appendix A.

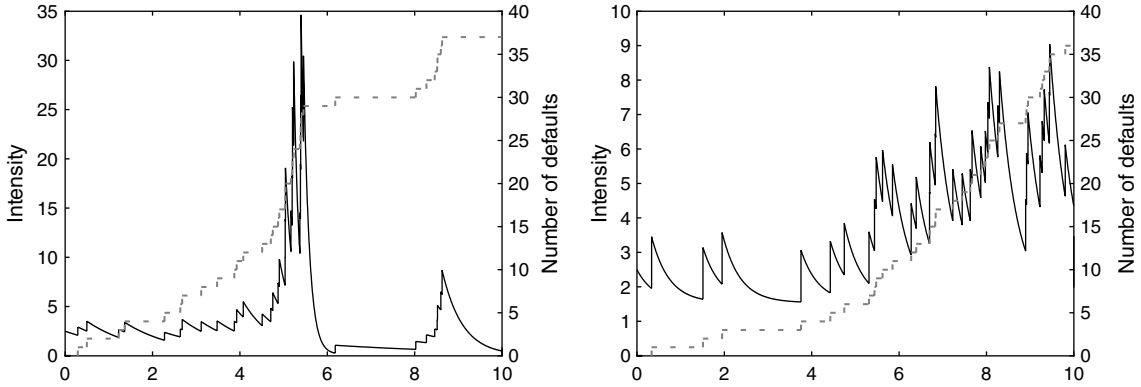
Proposition 3.1. *If $c(1 + \delta) < 1$, then the counting process N is nonexplosive.*

The intensity (10) follows a piecewise deterministic process with right-continuous sample paths. It jumps at an event time T_n. The jump magnitude is random. It is equal to $\max(\gamma, \delta\lambda_{T_n^-})$ and depends on the intensity just before the event, which itself is a function of the event times $T_1, \ldots, T_{n-1}$. The minimum jump size is γ. From T_n onward the intensity reverts exponentially to the level $c\lambda_{T_n}$, at rate $\kappa\lambda_{T_n}$. Because the reversion rate and level are proportional to the value of the intensity at the previous event, they depend on the times $T_1, \ldots, T_n$ and change adaptively at each default. For $T_n \leqslant t < T_{n+1}$, the behavior of the intensity is described by the $\mathcal{F}_{T_n}$-measurable function

$$\lambda_t = c\lambda_{T_n} + (1 - c)\lambda_{T_n} \exp(-\kappa\lambda_{T_n}(t - T_n)). \tag{12}$$

The dependence of the reversion level c_t, reversion speed κ_t, and jump magnitude $\max(\gamma, \delta\lambda_{T_n^-})$ on the path of the counting process N distinguishes our specification (10) from the classical (Hawkes 1971) model. The Hawkes intensity follows a piecewise deterministic process $d\lambda_t = \kappa(c - \lambda_t)\,dt + \delta\,dU_t$, where $U = u_1 + \cdots + u_N$ and the jump

Figure 2. Left panel: Sample path of the intensity (left scale, solid) and default process (right scale, dotted) for the adaptive model (10) with $\theta = (0.25, 0.05, 0.4, 0.8, 2.5)$, values that are motivated by our estimation results in §3.4. Right panel: Sample path of the intensity and default process for the Hawkes model $d\lambda_t = \kappa(c - \lambda_t)\,dt + \delta\,dU_t$ where the jump magnitudes $u_n = 1$ so $U = N$, $\lambda_0 = 2.5$, and $\kappa = 2.5$ and $c = \delta = 1.5$ are chosen so that the expected number of events over 10 years matches that of the model (10), roughly 37.



Notes. The reversion level and speed change at each event. Algorithm 3 is used to generate the paths. Note that the Hawkes intensity cannot fall below the global reversion level c. The algorithm in Ogata (1981) is used to generate the Hawkes model paths.

magnitudes u_n are drawn from a fixed distribution on $\mathbb{R}_+$. This model is more rigid than our adaptive model (10): it imposes a global, state-independent reversion level c and speed κ, and the magnitude of a jump in the Hawkes intensity is drawn independently of past event times. Figure 2 contrasts the sample paths of (λ, N) for the Hawkes model with those for our model (10). The paths exhibit different clustering behavior, with the Hawkes model generating a more regular clustering pattern. While Azizpour et al. (2008a) found that a variant of the Hawkes model performs well on the economy-wide default data, we had difficulty fitting this and several other variants of the Hawkes model to the portfolio default times generated by the resampling mechanism. We found the constant reversion level and speed to be too restrictive. Our adaptive specification (10) relaxes these constraints while preserving parsimony and computational tractability.

The jumps of the intensity process (10) are statistically important. They generate event correlation: an event increases the likelihood of further events in the near future. This feature facilitates the replication of the event clusters seen in Figure 1, a fact that we establish more formally below. The intensity jumps can also be motivated in economic terms. They represent the impact of a default on the other firms, which is channeled through the complex web of contractual relationships in the economy. The existence of these feedback phenomena is indicated by the ripple effects associated with the default of Lehman Brothers on September 15, 2008, and is further empirically documented in Azizpour et al. (2008a), Jorion and Zhang (2007), and others. Our jump size specification guarantees that the impact of an event increases with the default rate prevailing at the event: the weaker the firms the stronger the impact. An alternative motivation of the intensity jumps is Bayesian learning: an event reveals information about the values of unobserved event covariates, and this leads to an update of the intensity; see Collin-Dufresne et al. (2009), Duffie et al. (2009), and others.

3.3. Event and Loss Simulation

The piecewise deterministic dynamics of the intensity (10) generate computational advantages for model calculation. One benefit of this feature is that it allows us to estimate the distribution of N and related quantities by *exact* Monte Carlo simulation of the jump times of N. The interarrival times of N can be generated sequentially by acceptance-rejection sampling from a dominating counting process with intensity $\lambda_{T_{N_{t-}}} \geqslant \lambda_{t-}$.

Algorithm 3 (Default time simulation)
Generating a sample path of the portfolio default process N over $[\tau, H]$ for a horizon $H > \tau$ and the model (10).

Draw $(N_\tau, T_{N_\tau}, \lambda_{T_{N_\tau}})$ uniformly from $\{(N_\tau(\omega_i), T_{N_\tau}(\omega_i), \lambda_{T_{N_\tau}}(\omega_i))\colon i \leqslant I\}$.
Initialize $(n, S) \leftarrow (N_\tau, \tau)$ and $\lambda_S \leftarrow c\lambda_{T_n} + (1-c)\lambda_{T_n}\exp(-\kappa\lambda_{T_n}(S - T_n))$.
loop
  Draw $\mathscr{E} \sim \mathrm{Exp}(\lambda_S)$ and set $T \leftarrow S + \mathscr{E}$.
  if $T > H$ **then**
    Exit loop.
  end if
  $\lambda_T \leftarrow c\lambda_{T_n} + (\lambda_S - c\lambda_{T_n})\exp(-\kappa\lambda_{T_n}(T - S))$
  Draw $u \sim U(0, 1)$.
  if $u \leqslant \lambda_T/\lambda_S$ **then**
    $\lambda_T \leftarrow \lambda_T + \max(\gamma, \delta\lambda_T)$
    Assign $T_{n+1} \leftarrow T$ and update $n \leftarrow n + 1$.
  end if

14: Set $S \leftarrow T$ and $\lambda_S \leftarrow \lambda_T$.
15: **end loop**

The steps are summarized in Algorithm 3. In Step 1, we draw the state at the analysis time τ from the resampling scenarios of N and the corresponding fitted intensity scenarios. The rate λ_S of the dominating Poisson process, from which a candidate default time is generated in Step 4, is not only redefined at each acceptance but also at each rejection of a candidate time. This improves the efficiency of the algorithm. The acceptance probability is increased, and fewer candidate times are wasted.

Algorithm 3 leads to a trajectory of the default process N for a portfolio with replacement of defaulters. A trajectory of the corresponding portfolio loss process $L = l_1 + \cdots + l_N$ is obtained by drawing the loss l_n at each event time T_n of N. The distribution μ_l of l_n is estimated by the empirical distribution of the losses associated with all defaults in the resampling scenarios $\{N(\omega_i)\colon i \leqslant I\}$.

Once a path of N is generated, Algorithm 2 can be used to obtain a path of the default process N' for a portfolio without replacement of defaulters. A trajectory of the corresponding portfolio loss process $L' = l'_1 + \cdots + l'_{N'}$ is obtained by drawing the loss l'_n at each event time from μ_l.

Algorithm 3 is easy to implement and runs fast. For example, generating 100 K paths of N over five years for the fitted model parameters in Table 1 takes 1.67 seconds. Generating 100 K paths of N' takes, in the same configuration, 2.14 seconds.[9]

3.4. Likelihood Estimators

Another advantage of the piecewise deterministic intensity dynamics (10) is that the log-likelihood function in problem (5) takes a closed form that can be computed exactly. Based on $I = 10$ K resampling scenarios, we address problem (5) with the Nelder-Mead algorithm, which uses only function values. The algorithm is initialized at a set of random parameter values, which are drawn from a uniform distribution on the parameter space $\Theta = (0, 2) \times (0, 1) \times (0, 2)^2 \times (0, 20)$. For each of 100 randomly chosen initial parameter sets, the algorithm converges to the optimal parameter value $\hat{\theta}$ reported in Table 1.[10] We also provide asymptotic and bootstrapping confidence intervals. The left panel of Figure 3 shows the path of the fitted mean intensity $\hat{\lambda} = I^{-1} \sum_{i=1}^{I} \lambda^{\hat{\theta}}(\omega_i)$.

To provide some perspective on the parameter estimates, we employ an alternative inference procedure. Instead of maximizing the total likelihood (5) associated with all paths of N, we maximize the path log-likelihood

$$\sup_{\theta(\omega_i)\in\Theta} \int_0^\tau (\log \lambda_{s-}^{\theta(\omega_i)}(\omega_i)\, dN_s(\omega_i) - \lambda_s^{\theta(\omega_i)}(\omega_i)\, ds)$$

for each $i = 1, 2, \ldots, I$. The last row in Table 1 shows the median of the empirical distribution of the per-path MLE $\theta(\omega_i)$ over all paths ω_i. These values are in good agreement with the MLEs, supporting our total likelihood estimation strategy.

3.5. Testing In-Sample Fit

Above we have stressed the computational advantages of our intensity model (10). In this section and the next, we evaluate the model statistically. We show that the model fits the default data, and that the fitted model leads to accurate event forecasts.

We evaluate the joint specification of Z^* and λ based on Proposition 2.2.[11] The right panel of Figure 3 shows the fitted path of the economy-wide intensity $\lambda^* = \lambda / Z^*$, which defines the time change. We test the Poisson property of the time-scaled interarrival times $W_n^* = \int_{T_{n-1}^*}^{T_n^*} 1/Z_s^* \lambda_s\, ds$ using a Kolmogorov-Smirnov (KS) test, a QQ plot, and Prahl's (1999) test. The KS test addresses the deviation of the empirical distribution of the (W_n^*) from their theoretical standard exponential distribution. Prahl's test is particularly sensitive to large deviations of the (W_n^*) from their theoretical mean 1. Prahl shows that if the (W_n^*) are independent samples from a standard exponential distribution, then

$$M = \frac{1}{m} \sum_{n:\, W_n^* < \mu} \left(1 - \frac{W_n^*}{\mu}\right) \tag{13}$$

is asymptotically normally distributed with mean $\mu_M = e^{-1} - 0.189/m$ and standard deviation $\sigma_M = 0.2427/\sqrt{m}$, where $m = N_\tau^*$ is the number of arrivals and μ is the sample mean of $(W_n^*)_{n \leqslant m}$. We reject the null hypothesis of a correct joint specification of Z^* and λ if the test statistic $\Delta_M = (M - \mu_M)/\sigma_M$ lies outside of the interval $(-1, 1)$.

Table 1. Maximum likelihood estimates of the parameters of the intensity λ for the CDX.HY6 portfolio, along with estimates of asymptotic (A) and bootstrapping (B) 95% confidence intervals (10 K bootstrap samples were used).

Parameter	κ	c	δ	γ	λ_0
MLE	0.254	0.004	0.419	0.810	8.709
95% CI (A)	(0.252, 0.255)	(0.004, 0.004)	(0.417, 0.422)	(0.806, 0.813)	(8.566, 8.851)
95% CI (B)	(0.250, 0.258)	(0.003, 0.004)	(0.406, 0.433)	(0.791, 0.829)	(8.538, 8.882)
Median	0.263	0.004	0.430	0.779	9.405

Notes. The "Median" row indicates the median of the empirical distribution of the per-path MLEs over all resampling paths. The estimates are based on $I = 10$ K resampling scenarios, generated by Algorithm 1 from the observed defaults in the universe of Moody's rated names from 1/1/1970 to 11/7/2008.

Figure 3. Left panel: Fitted mean portfolio intensity $\hat{\lambda} = I^{-1}\sum_{i=1}^{I} \lambda^{\hat{\theta}}(\omega_i)$ vs. [5%, 95%] percentiles (boxes), [1%, 99%] percentiles (whiskers) and the mean number of portfolio defaults in any given year between 1970 and 2008, based on $I = 10$ K resampling scenarios for the CDX.HY6. Right panel: Fitted economy-wide intensity $\hat{\lambda}/Z^*$ vs. economy-wide defaults between 1970 and 2008, semiannually. The fitted intensity matches the time-series fluctuation of observed economy-wide default rates.

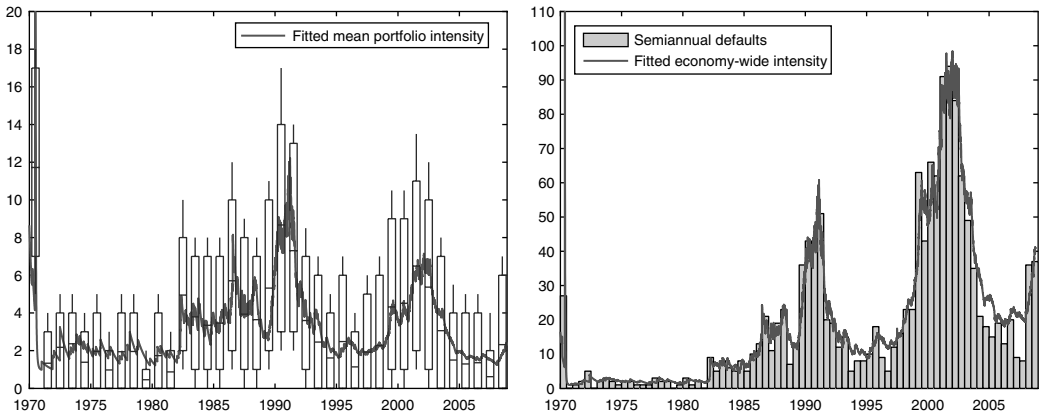


Figure 4 contrasts the empirical distribution of the (W_n^*) with the theoretical standard exponential distribution. The KS test has a p-value of 0.094, indicating that the deviation of the empirical distribution from the theoretical distribution is not statistically significant at standard confidence levels. The value of Prahl's test statistic $\Delta_M = 0.87$ leads to the same conclusion. The results of these and other[12] tests suggest that the economy-wide intensity λ^* generated by (3) and (10) replicates the substantial time-series variation of default rates during 1970–2008. This means that our adaptive intensity model (10) captures the clustering of defaults observed during the sample period.

A benefit of the time-scaling tests is that they can be applied to alternative model formulations, facilitating a direct comparison of fitting performance. Consider the bottom-up specification of Das et al. (2007), which appears to be much richer than ours: they specify firm-level intensity models with conditioning information given by a set of firm-specific and macroeconomic covariates. Das et al. (2007) find that the time-scaled, economy-wide times generated by their model deviate significantly from those of a Poisson process. In particular, they find that their specification does not completely capture the event clusters in the data. Thus, they reject this model formulation at standard confidence levels. This may indicate that, relative to our portfolio-level formulation with conditioning information given by past default timing, the additional modeling and estimation effort involved in a firm-level intensity model formulation with a large set of exogenous covariates might not translate into better fits and forecast performance.

Figure 4. In-sample fitness tests: empirical distribution of time-scaled, economy-wide interarrival times generated by the fitted (mean) paths of Z^* and λ, based on $I = 10$ K resampling scenarios for the CDX.HY6. Left panel: Empirical quantiles of time-scaled interarrival times vs. theoretical standard exponential quantiles. Right panel: Empirical distribution function (solid) of time-scaled interarrival times vs. theoretical standard exponential distribution function (dotted) along with 1% and 5% bands.

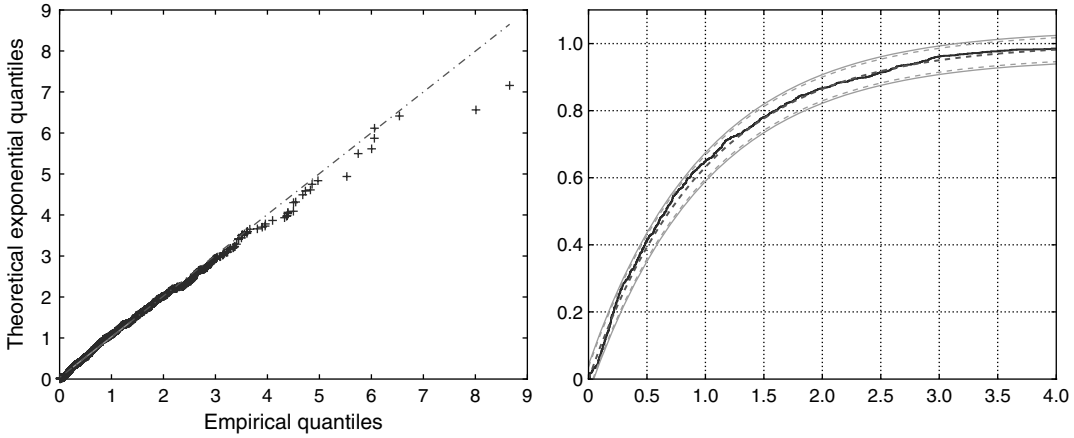
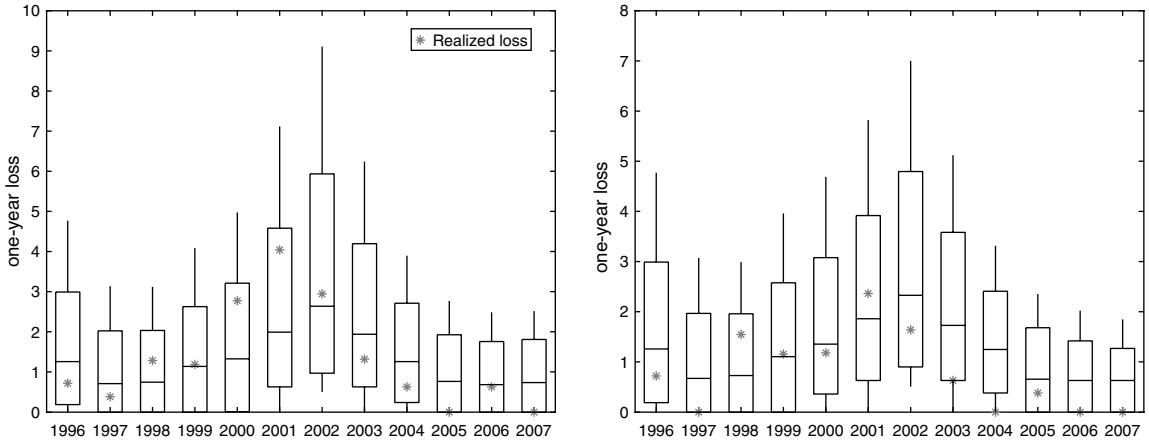


Figure 5. Out-of-sample test of loss forecasts for a test portfolio with the same rating composition as the CDX.HY6. Left panel: The test portfolio is selected at the beginning of each period. Right panel: The test portfolio is selected in 1996.



Note. We show the [25%, 75%] percentiles (box), the [15%, 85%] percentiles (whiskers) and the median (horizontal line) of the fitted conditional distribution of the incremental portfolio loss $L'_{\tau+1} - L'_\tau$ given $\mathcal{F}_\tau$ for τ varying annually between 1/1/1996 and 1/1/2007, estimated from 100 K paths generated by Algorithms 2 and 3, and the realized portfolio loss during $[\tau, \tau+1]$.

Finally, we note that the time-scaling tests can also be used to evaluate the specification of λ directly on the resampling scenarios $\{N(\omega_i)\colon i \leqslant I\}$. For each path ω_i, we time-scale the portfolio default process $N(\omega_i)$ with its fitted compensator $A(\omega_i) = \int_0^{\cdot} \lambda_s(\omega_i)\, ds$, and calculate the statistics of the KS test and Prahl's test for the time-scaled interarrival times $W_n = \int_{T_{n-1}(\omega_i)}^{T_n(\omega_i)} \hat{\lambda}_s^{\hat{\theta}}(\omega_i)\, ds$, which can be calculated exactly thanks to formula (12). The 95% confidence interval for the p-value of the KS test is 0.24 ± 0.01. The 95% confidence interval for Prahl's test statistic, defined for (W_n) in analogy to (13), is 0.81 ± 0.07. Also, these tests indicate that the adaptive model (10) is well specified.

3.6. Testing Out-of-Sample Loss Forecasts

Next we assess the out-of-sample event forecast performance of the fitted model (10). We contrast the loss forecast for a test portfolio without replacement with the actual loss in the test portfolio. The firms in the test portfolio are randomly selected from the universe of rated issuers, such that the test portfolio has the same rating composition as the CDX.HY6.[13] We apply Algorithms 2 and 3 to obtain an unbiased estimate of the conditional distribution of the incremental test portfolio loss $L'_{\tau+1} - L'_\tau$ given $\mathcal{F}_\tau$. We then contrast this distribution with the actual loss in the test portfolio during $(\tau, \tau+1]$, for yearly analysis times τ between 1/1/1996 and 1/1/2007.

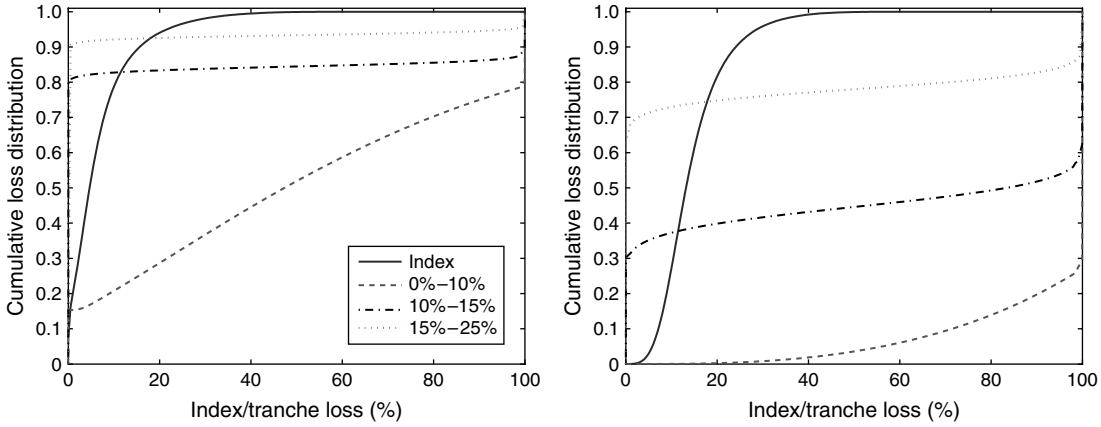
Figure 5 shows the results for two settings that represent common situations in practice. In one setting (left panel), we roll over the test portfolio in each one-year test period. That is, for each test period, we select a new test portfolio of 100 issuers at the beginning of the period, estimate the portfolio loss distribution based on data available at the beginning of the period, and compare it with the realized portfolio loss during the forecast period. In the other setting (right panel), we select the test portfolio only once, in 1996. For each period, we estimate the loss distribution for the portfolio of names that have survived to the beginning of the period, and we compare it with the realized portfolio loss during the forecast period. This setting is motivated by the situation of a buy and hold investor. The graphs indicate that the portfolio loss forecasts are very accurate.

4. Synthetic Collateralized Debt Obligations

Having established the fitting and forecast performance of the adaptive model (10) of portfolio default timing, we use it to estimate the exposure of an investor selling default protection on a tranche of a synthetic CDO. A synthetic CDO is based on a portfolio of C single-name credit swaps with notional 1 and maturity date H, the maturity date of the CDO. The constituent credit swaps are referenced on bonds issued by rated firms. When a bond issuer defaults, the corresponding credit swap pays the loss associated with the event. A defaulter is not replaced.

A tranche of a synthetic CDO is a swap contract specified by a lower attachment point $\underline{K} \in [0, 1)$, and an upper attachment point $\overline{K} \in (\underline{K}, 1]$. An index swap has attachment points $\underline{K} = 0$ and $\overline{K} = 1$. The protection seller agrees to cover all losses due to default in the reference portfolio, provided these losses exceed a fraction $\underline{K}$ of the total notional C of the reference contracts but are not larger than a fraction $\overline{K}$ of C. In exchange, the protection buyer pays to the protection seller an upfront fee at inception, and a quarterly fee, both of which are negotiated at contract inception.

Figure 6. Kernel smoothed conditional distribution of the normalized five-year cumulative tranche loss $U_H(\underline{K}, \overline{K})/C(\overline{K} - \underline{K})$ for the CDX.HY6 on 3/27/2006, the Series 6 contract inception date, for each of several standard attachment point pairs $(\underline{K}, \overline{K})$. Left panel: Actual distribution, estimated using the model and fitting methodology developed in §§2 and 3, based on $I = 10$ K resampling scenarios for N, and 100 K replications. Right panel: Risk-neutral distribution, estimated from the market prices paid for five-year tranche protection on 3/27/2006 using the method explained in online Appendix B, based on 100 K replications.



Note. Here, H represents the maturity date 6/27/2011.

We estimate the risk exposure of the tranche protection seller, which is measured by the conditional distribution at contract inception of the *cumulative tranche loss*, i.e., the portfolio loss allocated to the tranche over the life of the contract. With the convention that the portfolio loss at contract inception is equal to zero, the cumulative tranche loss at a post-inception time $t \leqslant H$ is given by the "call spread."

$$U_t(\underline{K}, \overline{K}) = (L'_t - C\underline{K})^+ - (L'_t - C\overline{K})^+. \tag{14}$$

The left panel of Figure 6 shows the conditional distribution of the normalized five-year cumulative tranche loss $U_H(\underline{K}, \overline{K})/C(\overline{K} - \underline{K})$ for the CDX.HY6 on 3/27/2006, the Series 6 contract inception date,[14] for each of several standard attachment point pairs. The maturity date H for Series 6 contracts is 6/27/2011.[15] To estimate the tranche loss distribution, we generate default scenarios during 1/1/1970–3/27/2006 for the CDX portfolio from Moody's default history by the resampling Algorithm 1. Next we estimate the intensity model (10) from these scenarios, as described in §3.4, with parameter estimates reported in Table 2. Then we generate event times during the prediction interval 3/27/2006–6/27/2011 from the fitted intensity using Algorithm 3, which we thin using Algorithm 2 to obtain paths of portfolio default times and losses without replacement. The trajectories for $U(\underline{K}, \overline{K})$ thus obtained lead to an unbiased estimate of the desired distribution.

The loss distributions indicate the distinctions between the risk profiles of the different tranches. The equity tranche, which has attachment points 0 and 10%, carries the highest exposure among all tranches. For the equity protection seller, the probability of trivial losses is roughly 15%, and the probability of losing the full tranche notional is 20%. For the (10%, 15%)-mezzanine protection seller, the probabilities are roughly 80% and 10%, respectively. The mezzanine tranche is less risky because the equity tranche absorbs the first 10% of the total portfolio loss. For the (15%, 25%)-senior protection seller, the probabilities are roughly 90% and 5%, respectively.

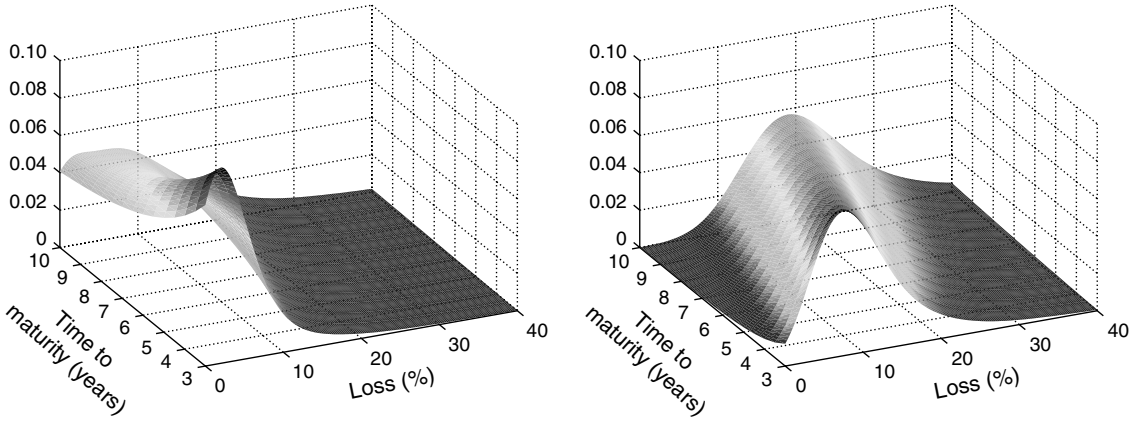
To provide some perspective on these numbers, we also estimate the *risk-neutral* tranche loss distributions implied

Table 2. Maximum likelihood estimates of the parameters of the portfolio intensity λ for the CDX.HY6, along with estimates of asymptotic (A) and bootstrapping (B) 95% confidence intervals (10 K bootstrap samples were used).

Parameter	κ	c	δ	γ	λ_0
MLE	0.260	0.003	0.439	0.856	8.494
95% CI (A)	(0.259, 0.261)	(0.003, 0.003)	(0.437, 0.441)	(0.853, 0.859)	(8.368, 8.621)
95% CI (B)	(0.256, 0.265)	(0.003, 0.004)	(0.424, 0.454)	(0.836, 0.876)	(8.327, 8.650)

Notes. The estimates are based on $I = 10$ K resampling scenarios for N, generated by Algorithm 1 from the observed economy-wide defaults in Moody's universe of rated names from 1/1/1970 to 3/27/2006.

Figure 7. Kernel smoothed conditional distribution of the normalized cumulative portfolio loss L'_H/C for the CDX.HY6 on 3/27/2006 for horizons H varying between 3/27/2009 and 3/27/2016. Left panel: Actual distribution, estimated using the model and fitting methodology developed in §§2 and 3, based on $I = 10$ K resampling scenarios for N, and 100 K replications. Right panel: Risk-neutral distribution, estimated from the market prices paid for five-year tranche protection on 3/27/2006 using the method explained in online Appendix B, based on 100 K replications.



by the market prices paid for five-year tranche protection on 3/27/2006. The estimation procedure, explained in online Appendix B, is based on the intensity model (10) relative to a risk-neutral pricing measure. The risk-neutral intensity parameters are chosen such that the model prices match the market index and tranche prices as closely as possible. The fitting errors are small, which indicates that our adaptive intensity model (10) performs also well under a risk-neutral pricing measure.

The right panel of Figure 6 shows the risk-neutral distribution of $U_H(\underline{K}, \overline{K})/C(\overline{K} - \underline{K})$ for the CDX.HY6, for each of several standard attachment point pairs. For the equity protection seller, the risk-neutral probability of trivial losses is zero, and the risk-neutral probability of losing the full tranche notional is 70%. Compare with the actual probabilities indicated in the left panel of Figure 6. For any tranche, the risk-neutral probability of losing more than any given fraction of the tranche notional is much higher than the corresponding actual probability. The distinction between the probabilities reflects the risk premium the protection seller requires for bearing exposure to the correlated corporate default risk in the reference portfolio. Figure 7 shows the distributions of the normalized cumulative portfolio loss L'_H/C for multiple horizons H. For all horizons, the risk-neutral distribution has a much fatter tail than the corresponding actual distribution. In other words, when pricing index and tranche contracts, the market overestimates the probability of extreme default scenarios relative to historical default experience.

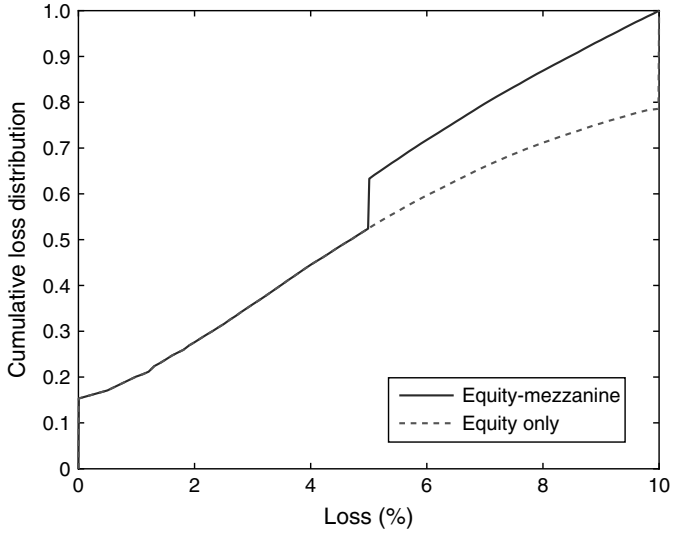
Our tools can also be used to analyze more complex investment positions. Investors often trade the CDO capital structure, i.e., they simultaneously sell and buy protection in different tranches. A popular trade used to be the "equity-mezzanine" trade, in which the investor sells equity protection and hedges the position by buying mezzanine protection with matching notional and maturity. Figure 8 contrasts the conditional distribution on 3/27/2006 of the normalized cumulative loss $(U_H(0, 0.1) - U_H(0.1, 0.15))/0.1C$ generated by this trade with the conditional distribution of the normalized cumulative equity loss $U_H(0, 0.1)/0.1C$, both for the CDX.HY6. While the mezzanine hedge does not alter the probability of trivial losses in the equity-only position, it does reduce the probability of a total loss of notional from 20% to virtually zero. This is because the mezzanine protection position generates cash flows when equity is wiped out. The magnitude of these cash flows is, however, capped at 50% of the total equity notional, and this property generates the point mass at 5% in the loss distribution of the equity-mezzanine position.

We can also measure the risk of positions in tranches referenced on different portfolios. For example, an investor may sell and buy protection on several of the reference portfolios in the CDX family, including the High Yield, Investment Grade and Crossover portfolios. In this case, we estimate default and loss processes for each of the portfolios based on the realization of the economy-wide default process. We generate events for each portfolio and then aggregate the corresponding position losses as in the equity-mezzanine case.

5. Cash Collateralized Debt Obligations

Next we use the adaptive model to estimate the exposure of an investor in a tranche of a cash CDO. A cash CDO is based on a portfolio of corporate bonds, mortgages, or other credit obligations, which is bought by a special purpose vehicle (SPV) that finances the purchase by issuing a collection of tranches that might pay coupons. The interest

Figure 8. Kernel smoothed conditional distribution on 3/27/2006 of the normalized cumulative loss $(U_H(0, 0.1) - U_H(0.1, 0.15))/0.1C$ associated with selling equity protection and buying mezzanine protection with matching notional and maturity, along with the distribution of the normalized cumulative equity loss $U_H(0, 0.1)/0.1C$, for the CDX.HY6.

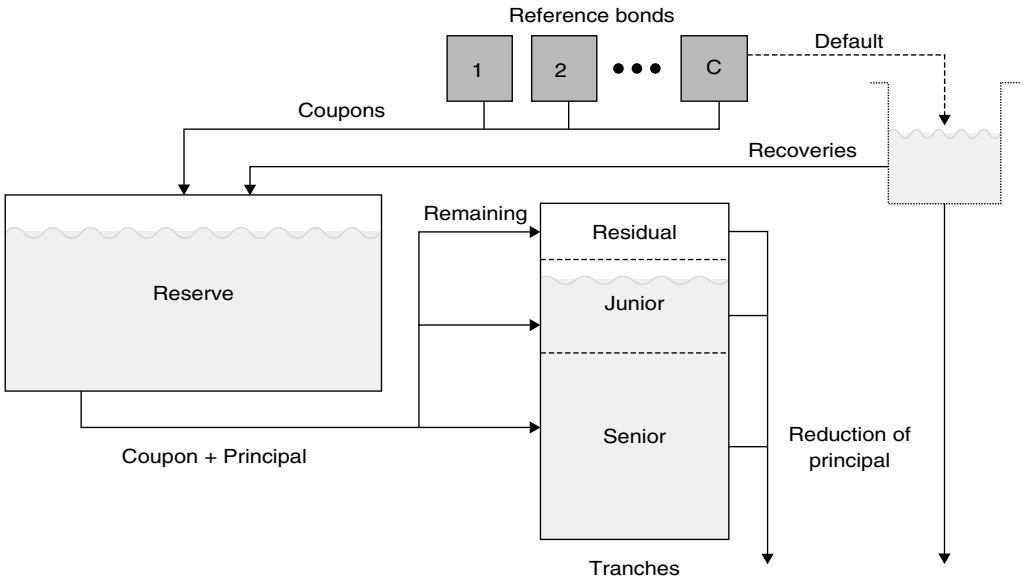


and recovery cash flows generated by the reference bonds are collected by the SPV in a reserve account and are then allocated to the tranches according to a specified prioritization scheme. Losses due to defaults are absorbed by the tranche investors, usually in reverse priority order. Figure 9 illustrates a typical structure. The reference portfolio of a cash CDO may be actively managed. In this case, the asset manager can buy and sell collateral bonds and replace defaulted obligations. We analyze a basic unmanaged cash CDO, in which the reference bonds are held to maturity and defaulted bonds are not replaced. We assume that the C reference bonds are straight coupon bonds with notional 1, maturity date H equal to the CDO maturity date, issuance date t_0 equal to the CDO inception date, coupon dates $(t_m)_{1\leqslant m\leqslant M}$ with $t_M = H$, and per-period coupon rate v. The total interest income to the SPV in period m is therefore

$$W(m) = v(C - N'_{t_m}). \tag{15}$$

When a reference bond defaults, the SPV collects the recovery at the coupon date following the event. The total

Figure 9. Cash flow "waterfall" of a sample cash CDO.



recovery cash flow in period m is

$$K(m) = N'_{t_m} - N'_{t_{m-1}} - (L'_{t_m} - L'_{t_{m-1}}).$$

The total cash flows from coupons and recoveries are invested in a reserve account that earns interest at the risk-free rate r. For period m, they are given by

$$B(m) = \begin{cases} W(m) + K(m) & \text{if } 1 \leqslant m < M, \\ W(m) + K(m) + C - N'_{t_m} & \text{if } m = M. \end{cases}$$

The SPV issues three tranches, of which two are debt tranches that promise to pay a specified coupon at times (t_m). There is one senior debt tranche, represented by a sinking-fund bond with initial principal p_1 and per-period coupon rate c_1, and a junior debt tranche that is represented by a sinking-fund bond with initial principal p_2 and per-period coupon rate c_2. The initial principal of the residual equity tranche is $p_3 = C - p_1 - p_2$. Each tranche has maturity date H.

For the debt tranches $j = 1, 2$ and coupon period m, we denote by $F_j(m)$ the remaining principal. The scheduled interest payment is $F_j(m)c_j$, and the accrued unpaid interest is $A_j(m-1)$. If $Q_j(m)$, the actual interest paid in period m, is less than $F_j(m)c_j$, then the difference is accrued at c_j to generate an accrued unpaid interest of

$$A_j(m) = (1 + c_j)A_j(m-1) + c_j F_j(m) - Q_j(m).$$

The actual payments to the debt tranches are prioritized. We consider two prioritization schemes, the uniform and fast schemes, which were introduced by Duffie and Garleanu (2001) and are reviewed in online Appendix C for completeness. A prioritization scheme specifies the actual interest payment $Q_j(m)$, the prepayment of principal $P_j(m)$, and the contractual unpaid reduction in principal $J_j(m)$. The total cash payment in period m is $Q_j(m) + P_j(m)$. The remaining principal after interest payments is

$$F_j(m) = F_j(m-1) - P_j(m) - J_j(m) \geqslant 0.$$

The *par coupon rate* on tranche j is the scheduled coupon rate c_j with the property that the initial market value of the bond is equal to its initial face value $F_j(0) = p_j$.

The residual equity tranche does not make scheduled coupon or principal payments, so $c_3 = 0$. Instead, at maturity H, after the debt tranches have been paid as scheduled, any remaining funds in the reserve account are allocated to the equity tranche.

We analyze the exposure of a tranche investor for a five-year cash CDO referenced on the CDX.HY6 of $C = 100$ names. This choice of reference portfolio allows us to compare the results with those for the synthetic CDO. Before we can start this analysis, we need to price the reference bonds and the debt tranches.[16] The first step is to estimate the par coupon rate of the reference bonds. This is the scheduled coupon rate v^* with the property that the initial market value of a reference bond is equal to its initial face value 1. Given v^*, the second step is to estimate the par coupon rates c_j^* of the debt tranches. Online Appendix D gives details on these steps, and Table 3 reports the annualized par coupon rates and spreads. The par coupon spread is the difference between the par coupon rate and the (hypothetical) par coupon rate that would be obtained if the reference bonds were not subject to default risk.

Next we estimate the conditional distribution of the discounted cumulative tranche loss, which is the loss the tranche investor faces during the life of the contract, discounted at the risk-free rate r. At a post-inception time $t \leqslant H$, the discounted cumulative tranche loss is given by

$$U_{jt}(H, p_j, v^*, c_1^*, c_2^*) = \overline{V}_{jt}(H, p_j, v^*, c_1^*, c_2^*) - V_{jt}(H, p_j, v^*, c_1^*, c_2^*), \tag{16}$$

where $V_{jt}(H, p_j, v^*, c_1^*, c_2^*)$ is the present value at time t of the coupon and principal cash flows actually paid to tranche j over the life of the tranche, and $\overline{V}_{jt}(H,\ p_j,\ v^*,\ c_1^*,\ c_2^*)$ is the present value of these cash flows assuming no defaults occur during the remaining life. For a debt tranche, $\overline{V}_{jt}(H, p_j, v^*, c_1^*, c_2^*)$ represents the present value of all scheduled coupon and principal payments. For the equity tranche, this quantity represents the present value of the reserve account value after the debt tranches have been paid as scheduled.

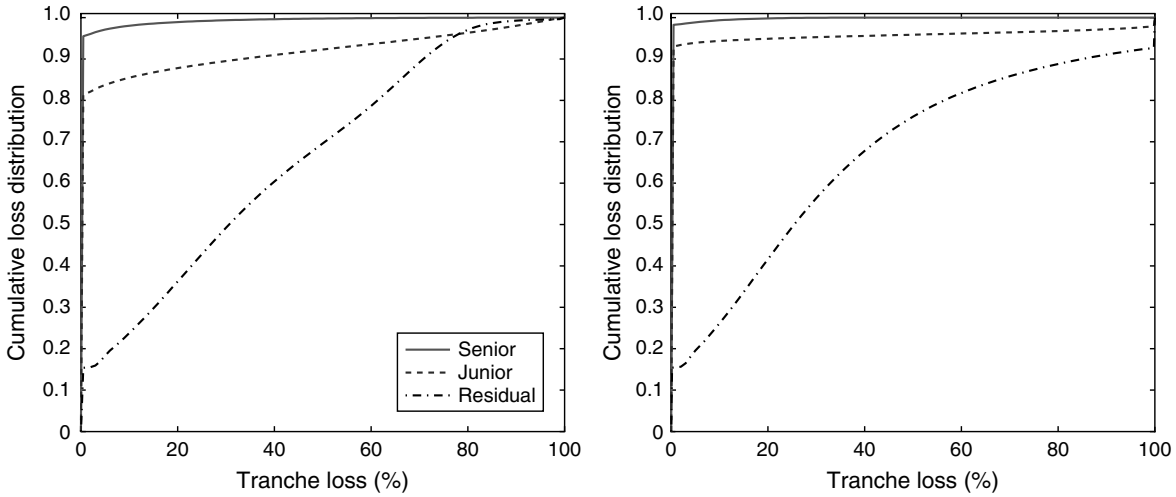
Note that due to the complex structure of the cash flow prioritization, the tranche loss $U_{jt}(H,\ p_j,\ v^*,\ c_1^*,\ c_2^*)$ does

Table 3. Fitted annualized par coupon rates and spreads on 3/27/2006 for the constituent bonds and debt tranches of a five-year cash CDO referenced on the CDX.HY6, for each of two standard prioritization schemes.

	Reference bond	Uniform scheme		Fast scheme	
		Senior	Junior	Senior	Junior
Annualized par coupon rate (bp)	811.23	476.58	1,299.42	483.38	749.75
Par coupon spread (bp)	339.20	23.68	846.52	11.29	277.77

Notes. The principal values are $(p_1, p_2, p_3) = (85, 10, 5)$. The fitting procedure is described in online Appendix D.

Figure 10. Kernel smoothed conditional distribution on 3/27/2006 of the normalized discounted loss $U_{j\tau}(H, p_j, v^*, c_1^*, c_2^*)/\overline{V}_{j\tau}(H, p_j, v^*, c_1^*, c_2^*)$ for a five-year cash CDO referenced on the CDX.HY6, whose maturity date H is 6/27/2011. Left panel: Uniform prioritization scheme. Right panel: Fast prioritization scheme.



Note. The initial tranche principals are $p_1 = 85$ for the senior tranche, $p_2 = 10$ for the junior tranche, and $p_3 = 5$ for the residual equity tranche. We apply the model and fitting methodology developed in §§2 and 3, based on $I = 10$ K resampling scenarios for N, and 100 K replications.

not admit a simple analytic expression in terms of the portfolio default count and loss of the reference portfolio. This leaves simulation as the only feasible tool for analyzing the cash CDO tranche loss, regardless of whether or not the models of the portfolio default count and loss processes are analytically tractable.

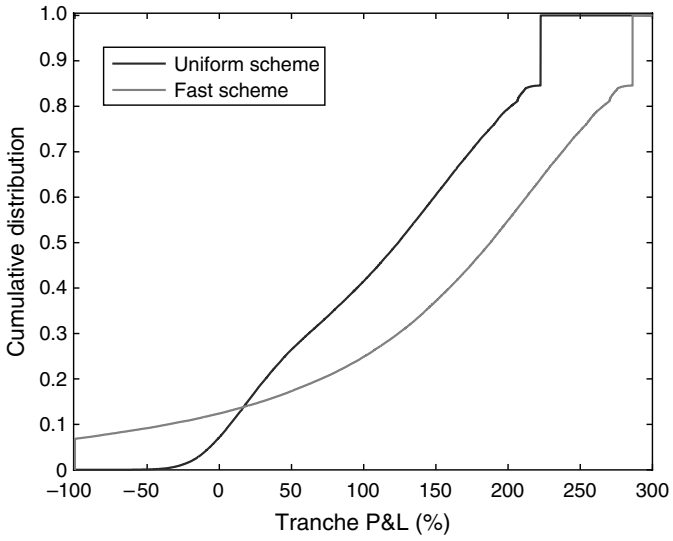
Figure 10 shows the fitted conditional distribution on 3/27/2006 of the normalized loss $U_{j\tau}(H, p_j, v^*, c_1^*, c_2^*)/\overline{V}_{j\tau}(H, p_j, v^*, c_1^*, c_2^*)$ for a five-year structure whose maturity date H is 6/27/2006, for each of several tranches. To estimate that distribution, we proceed as described in §4 for the synthetic CDO, and generate paths of the fitted portfolio default and loss processes. These are then fed into the cash flow calculator, which computes $V_{jt}(H, p_j, v^*, c_1^*, c_2^*)$ for each path. We assume that coupons are paid quarterly. The risk-free interest rate r is deterministic and is estimated from Treasury yields for multiple maturities on 3/27/2006, obtained from the website of the Department of Treasury.

As in the synthetic CDO case, the risk profile of a cash CDO tranche depends on the degree of overcollateralization, i.e., the location of the attachment points. It also depends on the prioritization scheme specified in the CDO terms. With the uniform scheme, the cash flows generated by the reference bonds are allocated sequentially to the debt tranches according to priority order, and default losses are applied to all tranches in reverse priority order. With the fast scheme, the cash flows are used to retire the senior tranche as soon as possible. After the senior tranche is retired, the cash flows are used to service the junior tranche. After the junior tranche is retired, any residual cash flows are distributed to the equity tranche. Therefore, the senior tranche investor is less exposed under the fast scheme: the probability of zero losses is increased from roughly 95% for the uniform scheme to roughly 98%. This reduction in risk is reflected by the par coupon spreads reported in Table 3. For the junior investor, the probability of zero losses increases from roughly 82% for the uniform scheme to roughly 93% for the fast scheme, but the probability of a loss equal to the present value of all scheduled coupon and principal payments increases from zero to about 3%. Nevertheless, the junior tranche commands a much smaller spread under the fast scheme. The risk reduction for the debt tranche investors comes at the expense of the equity investor: while the probability of zero losses is about 15% for both schemes, the probability of a loss equal to the present value of the reserve account value after the debt tranches have been paid as scheduled increases from zero to roughly 8% for the fast scheme. On the other hand, the equity investor has a much higher upside under the fast scheme. Figure 11 shows the distribution of the normalized discounted profit and loss $(V_{3\tau}(H, 5, v^*, c_1^*, c_2^*) - 5)/5$ for an equity tranche position. For the equity investor, the probability of more than doubling the principal p_3 after discounting is roughly 75% under the fast scheme, while it is only 60% under the uniform scheme. Figure 12 shows that for the debt tranches, the upside is much more limited, for any scheme. For example, the normalized discounted profit on the senior tranche can be at most 0.5% under the fast scheme and about 1% under the uniform scheme.

A standard measure to quantify the risk of a position is value at risk, a quantile of the position's loss distribution. We estimate the value at risk at time $t \leqslant H$ of a position in a cash CDO tranche maturing at H, given by

$$\mathrm{VaR}_{jt}(\alpha, H, p_j, v^*, c_1^*, c_2^*) = \inf\left\{x \in \mathbb{R}\colon P\left(\frac{U_{jt}(H, p_j, v^*, c_1^*, c_2^*)}{\overline{V}_{jt}(H, p_j, v^*, c_1^*, c_2^*)} \leqslant x \,\middle|\, \mathcal{F}_t\right) > \alpha\right\}$$

Figure 11. Kernel smoothed conditional distribution on 3/27/2006 of the normalized discounted profit and loss $(V_{3\tau}(H, 5, v^*, c_1^*, c_2^*) - 5)/5$ for the residual equity tranche of a five-year cash CDO referenced on the CDX.HY6, whose maturity date H is 6/27/2011. We apply the model and fitting methodology developed in §§2 and 3, based on $I = 10$ K resampling scenarios for N, and 100 K replications.

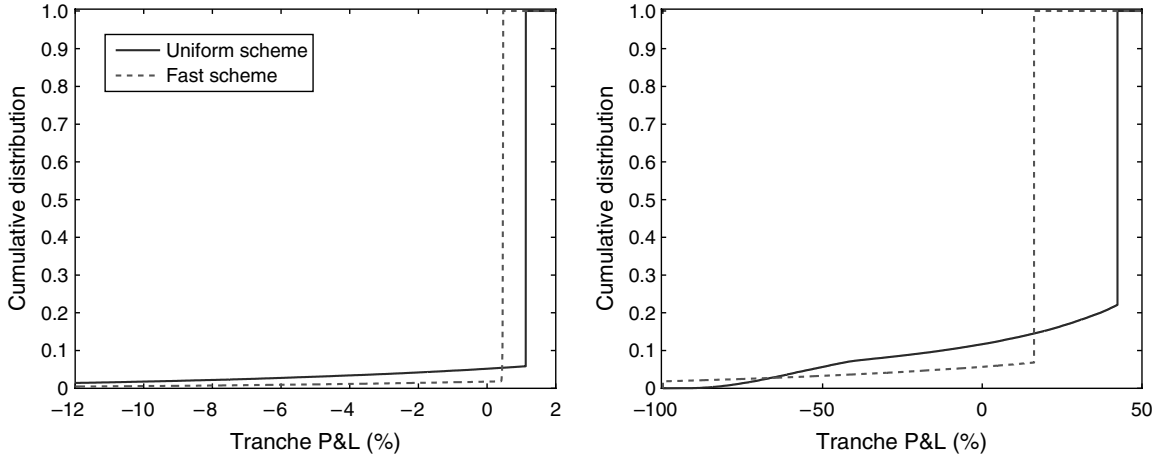


for some level of confidence $\alpha \in (0, 1)$. Figure 13 shows the 99.5% value at risk for the five-year senior tranche with maturity date H given by 6/27/2011, for analysis times varying weekly between 3/27/2006 and 11/7/2008, for each of the two prioritization schemes. Assuming risk capital is allocated according to the value at risk, a value in the time series represents the amount of risk capital that is needed at that time to support the tranche position over its remaining life. For both prioritization schemes, the time series behavior reflects the rising trend in corporate defaults that started in 2008 and that is evidenced in Figure 1. The fast scheme requires less risk capital than the uniform scheme but leads to higher volatility of risk capital.

6. Conclusion

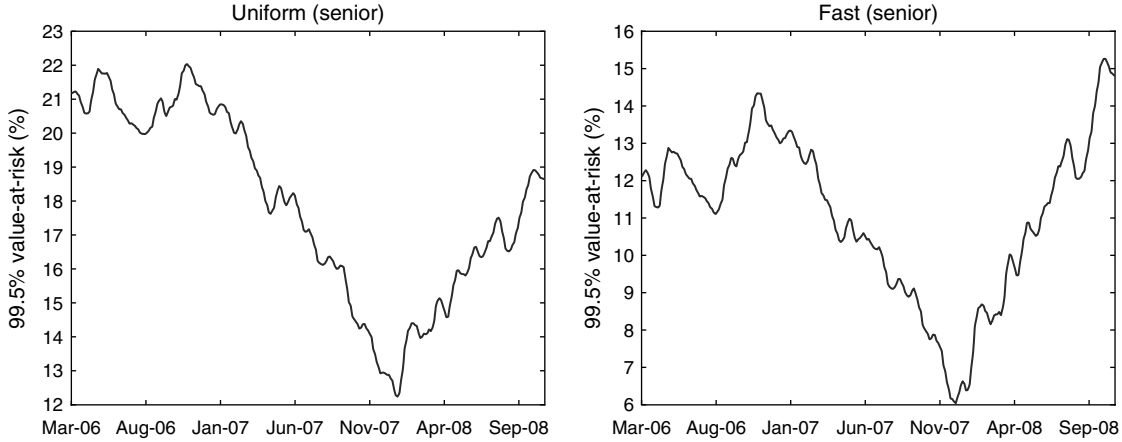
This paper develops, implements, and validates stochastic methods to measure the risk of investment positions in collateralized debt obligations and related credit derivatives tied to an underlying portfolio of defaultable assets.

Figure 12. Kernel smoothed conditional distribution on 3/27/2006 of the normalized discounted profit and loss $(V_{j\tau}(H, p_j, v^*, c_1^*, c_2^*) - p_j)/p_j$ for the debt tranches of a five-year cash CDO referenced on the CDX.HY6, whose maturity date H is 6/27/2011. Left panel: Senior tranche with principal $p_1 = 85$. Right panel: Junior tranche with principal $p_2 = 10$.



Note. We apply the model and fitting methodology developed in §§2 and 3, based on $I = 10$ K resampling scenarios for N, and 100 K replications.

Figure 13. 99.5% value at risk $\mathrm{VaR}_{1_\tau}(0.995, H, 85, v^*, c_1^*, c_2^*)$ for the five-year senior tranche of a cash CDO referenced on the CDX.HY6 with maturity date 6/27/2011, for τ varying weekly between 3/27/2006 and 11/7/2008, based on $I = 10$ K resampling scenarios for N, and 100 K replications. Left panel: Uniform prioritization scheme. Right panel: Fast prioritization scheme.



The ongoing financial crisis highlights the need for sophisticated yet practical tools allowing potential and existing investors as well as regulators to quantify the exposure associated with such positions, and to accurately estimate the amount of risk capital required to support a position.

The key to address the risk analysis problem is a model of default timing that captures the default clustering in the underlying portfolio and a method to estimate the model parameters based on historical default experience. This paper contributes to each of these two subproblems. It formulates an adaptive, intensity-based point process model of default timing that performs well according to in- and out-of-sample tests. Moreover, it develops a maximum likelihood approach to estimating point process models. This approach is based on an acceptance/rejection resampling scheme that generates alternative portfolio default histories from the available economy-wide default process.

The point process model and inference method have potential applications in other areas dealing with correlated event arrivals, within and beyond financial engineering. Financial engineering examples include the pricing and hedging of securities exposed to correlated default risk and order book modeling. Other example areas are insurance, health care, queuing, and reliability.

7. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://or.journal.informs.org/.

Endnotes

1. See Arnsdorf and Halperin (2008), Brigo et al. (2007), Chen and Glasserman (2008), Cont and Minca (2008), Ding et al. (2009), Duffie and Garleanu (2001), Eckner (2009), Errais et al. (2010), Kou and Peng (2009), Longstaff and Rajan (2008), Lopatin and Misirpashaev (2008), Mortensen (2006), Papageorgiou and Sircar (2007), and many others.
2. The limit in (1) exists and is equal to $Z_t^*(\omega)$ almost surely with respect to a certain measure on the product space $\Omega \times (0, \infty]$. See Giesecke et al. (2011) for more details.
3. Here and below, if Y is a right-continuous process with left limits, $Y_{t-} = \lim_{s \uparrow t} Y_s$.
4. Taking Z^* to be the predictable projection of the process given by formula (2) guarantees that Z^* is defined up to indistinguishability; see Dellacherie and Meyer (1982).
5. This choice makes (4) the so-called expected likelihood estimator. A sensitivity analysis indicates that the model parameter estimates reported in Table 1 are robust with respect to variations of α.
6. The asymptotic properties of the maximum likelihood estimator of θ are developed by Ogata (1978). The estimator is shown to be consistent, asymptotically normal, and efficient.
7. We distinguish issuers in terms of their "senior rating," which is an issuer-level rating generated by Moody's from ratings of particular debt obligations using its Senior Rating Algorithm (Hamilton 2005). We follow a common convention and subsume the categories Caa and Ca into the category C. We also subsume the numerical sub-categories Aa1, Aa2, Aa3 into the category Aa, and similarly for the other numerical sub-categories. Then the set of rating categories is $R = \{\text{Aaa, Aa, A, Baa, Ba, B, C, WR}\}$. The category WR indicates a withdrawn rating.
8. Each default event in the database has a time stamp; the resolution is one day. There are days with multiple events whose exact timing during the day cannot be established. Thus, the sequence of raw event dates in the data set is not strictly increasing. Algorithm 1 requires distinct event times, however. To address this issue, we assume that an event date is measured with uniformly distributed noise.

The noise is sampled when the paths of N are generated. We convert a raw economy-wide event date to a real-valued calendar time equal to 12 A.M. on the day of the event, and draw the noise from a uniform distribution on $[0, 1/365]$. With this choice, the randomization does not alter the original time stamp of an observed event. We have experimented with several alternative randomization schemes but have found that the model estimation results are insensitive to the chosen randomization scheme. Furthermore, since the data set allows us to measure the number of ρ-rated firms in the economy $X_t^*(\rho)$ only daily, in formula (3) for Z_t^* we take $X_{t-}^*(\rho)$ as the number of ρ-rated firms on the day prior to the day that contains t.

9. The numerical 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 the compiler used was Microsoft Visual C++ .NET version 7.1.3088. For random number generation, numerical optimization, and numerical root-finding, we used the GNU Scientific Library, Version 1.11.

10. In the computing environment described in footnote 9, it takes 0.15 seconds to evaluate the log-likelihood function for $I = 10$ K and a given parameter set. The full optimization takes 146.88 seconds when the initial parameters are set to the mid-points of the parameter space.

11. Proposition 2.2 requires the process Z^* to be strictly positive. The fitted values Z_t^* are indeed strictly positive during our sample period. We approximate the paths of the fitted process Z^* and the fitted mean intensity $\hat{\lambda}$ on a discrete-time grid with daily spacing. When multiple economy-wide defaults occur on the same day, then the arrival times are spaced equally.

12. To detect a possible serial dependence of the W_n^*, we also consider the test statistics SC1 and SC2 described by Lando and Nielsen (2009). These tests check the dependence of the number of re-scaled events in nonoverlapping time bins. Under the null of Poisson arrivals, the event counts are independent. We cannot reject the null for bin sizes 2, 4, 6, 8, and 10 at standard confidence levels.

13. We do not use the CDX.HY6 itself, because it was formed only in 2006, and we want to test the loss process over a longer time period, starting in 1996. The analysis would suffer from survivorship bias if we were to take the CDX portfolio itself.

14. Every six months, the CDX High Yield index portfolio is "rolled." That is, a new portfolio with a new serial number is formed by replacing names that have defaulted since the last roll, and possibly other names. The index and tranche swaps we consider are tied to a fixed series.

15. Although the actual time to maturity is five years and three months at contract inception, here and below, we follow the market convention of referring to the contract as a "five-year contract." At an index roll, new "five-year contracts" are issued, and these mature in five years and three months from the roll date.

16. In practice, this is done by the SPV, which then offers the tranche terms to potential investors. We have to perform this task here because we do not have access to actual data for the structure we analyze.

Acknowledgments

The authors are grateful for default data from Moody's and market data from Morgan Stanley. They thank Eymen Errais, Igor Halperin, and Jack Kim for helpful discussions. They are also grateful to three anonymous referees, an associate editor, Shahriar Azizpour, Lisa Goldberg, Steve Kou, and especially Jeremy Staum for insightful comments.

References

Arnsdorf, M., I. Halperin. 2008. BSLP: Markovian bivariate spread-loss model for portfolio credit derivatives. *J. Comput. Finance* **12** 77–100.

Azizpour, S., K. Giesecke, B. Kim. 2008a. Premia for correlated default risk. Working paper, Stanford University, Stanford, CA.

Azizpour, S., K. Giesecke, G. Schwenkler. 2008b. Exploring the sources of default clustering. Working paper, Stanford University, Stanford, CA.

Box, G., G. Tiao. 1992. *Bayesian Inference in Statistical Analysis*. Wiley, New York.

Brigo, D., A. Pallavicini, R. Torresetti. 2007. Calibration of CDO tranches with the dynamical generalized-Poisson loss model. *Risk* **20**(5) 70–75.

Chen, Z. H., P. Glasserman. 2008. Fast pricing of basket default swaps. *Oper. Res.* **56**(2) 286–303.

Collin-Dufresne, P., R. Goldstein, J. Helwege. 2009. How large can jump-to-default risk premia be? Modeling contagion via the updating of beliefs. Working paper, Carnegie Mellon University, Pittsburgh.

Cont, R., A. Minca. 2008. Recovering portfolio default intensities implied by CDO quotes. Working paper, Columbia University, New York.

Das, S., D. Duffie, N. Kapadia, L. Saita. 2007. Common failings: How corporate defaults are correlated. *J. Finance* **62** 93–117.

Dellacherie, C., P.-A. Meyer. 1982. *Probabilities and Potential*. North Holland, Amsterdam.

Delloye, M., J.-D. Fermanian, M. Sbai. 2006. Dynamic frailties and credit portfolio modeling. *Risk* **19**(1) 101–109.

Ding, X., K. Giesecke, P. Tomecek. 2009. Time-changed birth processes and multiname credit derivatives. *Oper. Res.* **57**(4) 990–1005.

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

Duffie, D., A. Eckner, G. Horel, L. Saita. 2009. Frailty correlated default. *J. Finance* **64** 2089–2123.

Eckner, A. 2009. Computational techniques for basic affine models of portfolio credit risk. *J. Comput. Finance* **15** 63–97.

Errais, E., K. Giesecke, L. Goldberg. 2010. Affine point processes and portfolio credit risk. *SIAM J. Financial Math.* **1** 642–665.

Giesecke, K., L. R. Goldberg, X. Ding. 2011. A top-down approach to multiname credit. *Oper. Res.* Forthcoming.

Glasserman, P., N. Merener. 2003. Numerical solution of jump-diffusion LIBOR market models. *Finance Stochastics* **7** 1–27.

Hamilton, D. 2005. Moody's senior ratings algorithm and estimated senior ratings. Moody's Investors Service.

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

Jorion, P., G. Zhang. 2007. Good and bad credit contagion: Evidence from credit default swaps. *J. Financial Econom.* **84**(3) 860–883.

Kou, S., X. Peng. 2009. Default clustering and valuation of collateralized debt obligations. Working paper, Columbia University, New York.

Lando, D., M. S. Nielsen. 2009. Correlation in corporate defaults: Contagion or conditional independence? Working paper, Copenhagen Business School.

Lewis, P., G. Shedler. 1979. Simulation of nonhomogeneous Poisson processes by thinning. *Naval Logist. Quart.* **26** 403–413.

Longstaff, F., A. Rajan. 2008. An empirical analysis of collateralized debt obligations. *J. Finance* **63**(2) 529–563.

Lopatin, A., T. Misirpashaev. 2008. Two-dimensional Markovian model for dynamics of aggregate credit loss. *Adv. Econometrics* **22** 243–274.

Mortensen, A. 2006. Semi-analytical valuation of basket credit derivatives in intensity-based models. *J. Derivatives* **13** 8–26.

Ogata, Y. 1978. The asymptotic behavior of maximum likelihood estimators of stationary point processes. *Ann. Inst. Statist. Math.* **30**(A) 243–261.

Ogata, Y. 1981. On Lewis' simulation method for point processes. *IEEE Trans. Inform. Theory* **27** 23–31.

Papageorgiou, E., R. Sircar. 2007. Multiscale intensity models and name grouping for valuation of multi-name credit derivatives. *Appl. Math. Finance* **15**(1) 73–105.

Prahl, J. 1999. A fast unbinned test of event clustering in Poisson processes. Working paper, Universität Hamburg, Hamburg, Germany.

SEC. 2008. Summary report of issues identified in the commission staff's examinations of select credit rating agencies. Report, U.S. Securities and Exchange Commission, Washington, DC.