Default Clustering and Valuation of Collateralized Debt Obligations*

Xianhua Peng and Steven Kou†
Columbia University

January 2009

Abstract

The recent financial turmoil has witnessed the powerful impact of the default clustering effect (i.e., one default event tends to trigger more default events in the future and cross-sectionally), especially on the market of collateralized debt obligations (CDOs). We propose a model based on *cumulative* default intensities that can incorporate the default clustering effect. Furthermore, the model is tractable enough to provide a direct link between single-name credit securities, such as credit default swaps (CDS), and multi-name credit securities, such as CDOs. The result of calibration to the recent market data, when Bear Sterns, Lehman Brothers, etc. collapsed and default correlation among firms was substantially high, shows that the model is promising.

Keywords: CDO; Collateralized Debt Obligation; Default clustering; Default correlation; Doubly stochastic

JEL classification: C15; C51; G33

*We would like to thank Zaiwen Wen for his helpful comments. We are also grateful to many people who offered insights into this work including the seminar participants at Columbia University, Georgia Institute of Technology, the conference participants at INFORMS annual meeting, 2008, and those at the Winter Simulation Conference, 2008.

†Corresponding author. Address: 312 S. W. Mudd Building, 500 West 120th Street, New York, NY 10027, USA. Phone: +1 212-854-4334. Fax: +1 212-854-8103. Email address: sk75@columbia.edu.

Default Clustering and Valuation of Collateralized Debt Obligations

Abstract

The recent financial turmoil has witnessed the powerful impact of the default clustering effect (i.e., one default event tends to trigger more default events in the future and cross-sectionally), especially on the market of collateralized debt obligations (CDOs). We propose a model based on *cumulative* default intensities that can incorporate the default clustering effect. Furthermore, the model is tractable enough to provide a direct link between single-name credit securities, such as credit default swaps (CDS), and multi-name credit securities, such as CDOs. The calibration result to the recent market data, when Bear Sterns, Lehman Brothers, etc. collapsed and default correlation among firms was substantially high, shows that the model is promising.

Keywords: CDO; Collateralized Debt Obligation; Default clustering; Default correlation; Doubly stochastic

JEL classification: C15; C51; G33

1 Introduction

The recent financial crisis, which led to the collapse of some major financial institutions such as Bear Sterns, Fannie Mae, Freddie Mac, Merrill Lynch, Lehman Brothers, Washington Mutual, and Wachovia, has changed the financial landscape completely. It is too early to tell the full impact of the crisis, as the events are still unfolding. But in terms of research, two issues are clear: (1) We need better models to incorporate the default clustering effect, i.e., one default event tends to trigger more default events both across time and cross-sectionally. (2) We need better models for multi-name credit securities, such as collateralized debt obligations (CDOs), the mispricing of which contributed significantly to the current financial crisis. This paper attempts to study these two issues.

By extending the intensity based model of Duffie and Gârleanu (2001), we propose a credit risk model based on cumulative default intensities that can incorporate the default clustering effect. Furthermore, the model is tractable enough to provide a direct link between single-name credit securities, such as credit default swaps (CDS), and multi-name credit securities, such as CDOs. The calibration of our model to the market data of CDOs and CDS on March 14, 2008, when Bear Sterns was near bankruptcy, and the data on September 16, 2008, right after Lehman Brothers filed for bankruptcy protection, shows that the model is promising.

1.1 Background of CDOs

There are two types of credit derivatives: the single-name derivatives, whose payoffs depend on the credit events of a single reference name, and the portfolio credit derivatives, whose payoffs depend on the credit events of a portfolio of reference names. For example, a popular single-name credit derivative is the CDS, which is a bilateral contract between a default protection buyer and a protection seller. More precisely, the protection buyer pays a periodic premium to the protection seller, in return for the cover of loss if the reference name defaults.

The best-known portfolio credit derivative is perhaps the CDO. A CDO is a debt security that is secured by a portfolio of defaultable assets, such as bonds and CDS. A CDO partitions its underlying portfolio into tranches, each of which corresponds to a specific portion of the default losses of the portfolio. A typical CDO structure is shown in Fig. 1, where the portfolio underlying the CDO has 125 names and the CDO comprises 6 tranches. The investor of the first tranche takes the first 3% default losses. The investor of the second tranche begins to take losses when the cumulative portfolio losses exceed 3% and continues to do so until the cumulative losses reach 6%, and so on for other tranches. In return, the investors receive periodic coupon payments. The first tranche is called the equity tranche, which is the most risky and hence offers the highest coupon. The last tranche is called the senior tranche, which is the last to take default losses (see Duffie and Singleton, 2003, for a comprehensive discussion of CDOs).

1.1.1 Use of CDOs and the Impact of CDOs on the Current Financial Crisis

The major use of CDOs is that they provide a way to create new securities with higher credit quality out of a portfolio of securities with low credit quality. In particular, the senior tranche of a CDO created from a portfolio of non-investment grade bonds can have high credit ratings such as AA or AAA if properly structured. Indeed, as Duffie (2007) points out, CDOs provide more supply of debt instruments to specialized investors such as pension funds and insurance companies who can only invest in securities with high credit quality. Thus, by introducing CDOs more investors are able to invest in assets such as high-yield bonds and mortgages, which they were not be able to do before. We refer interested readers to Duffie (2007) for a comprehensive discussion of the costs and benefits of credit risk transfer for the efficiency and stability of the financial system.

Due to their capacity of credit enhancement, CDOs have had significant impact in the current financial crisis that started in 2007. Through credit derivatives, such as CDOs, the credit risk was transferred to third-party investors. Consequently, the current credit crisis

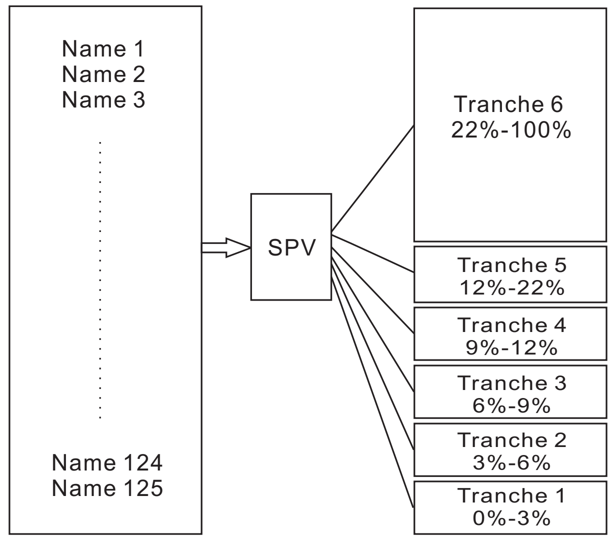

Figure 1: A typical CDO structure. SPV is the acronym of special purpose vehicle, which is usually a corporate entity that acquires a portfolio of assets and issues the CDO.

led to widespread losses to institutions and investors. Furthermore, the fear of the ubiquitous financial securities and the uncertainty in the financial markets caused lenders to reduce lending or lend at higher interest rates, which resulted in credit crunch and turmoils in the financial markets.

As mere instruments, CDOs themselves should not be held responsible for the subprime mortgage crisis. It is the people who misused them that should bear the main responsibility. There are many factors that have contributed to the subprime mortgage crisis (see Crouhy, Jarrow, and Turnbull, 2008, for a comprehensive examination). First, financial institutions that securitized those loans failed to perform due diligence. Demyanyk and Hemert (2008) find that even for six consecutive years before the crisis the quality of subprime mortgage loans deteriorated, but these financial institutions did not change their practice. Second, credit-rating firms have been criticized for giving over-optimistic ratings to subprime mortgage-backed securities and CDOs. The examination of the SEC uncovered weaknesses in their rating practices such as issues involving conflict of interests (SEC, 2008). Third, Duffie (2007) points out that the methodologies for rating CDOs are still at a relatively crude stage of development. Coval, Jurek, and Stafford (2008) find that credit ratings provided insufficient information for pricing and many investors who relied heavily on credit ratings for pricing could have overvalued the senior tranches of CDOs.

1.1.2 Difficulty of CDO Pricing

The main difficulty of CDO pricing comes from modeling the default correlation in the CDO portfolio. The marginal default probability of a single name in the portfolio can usually be implied from the market prices of certain securities such as bonds and CDS. The pricing of CDO tranches, however, requires the knowledge of the default correlation, i.e., the joint distribution of default times of underlying names, which cannot be implied easily from the market. Duffie (2007) writes that "even specialists in CDOs are currently ill equipped to measure the risks and fair valuation of tranches that are sensitive to default correlation. This is currently the weakest link in credit risk transfer markets."

1.2 Review of Existing CDO Pricing Models

In the rest of the paper, we assume that there are n names in the CDO portfolio and use τ_i to denote the default time of the i-th name, $i = 1, \ldots, n$. The two types of approaches in portfolio credit risk modeling are the bottom-up approach, which builds models for the correlated default times of individual names in the portfolio, and the top-down approach, which builds models for the cumulative loss of the whole portfolio without referring to the underlying single names.

1.2.1 Bottom-up Models

Bottom-up models can be separated into static models and dynamic models. Static models specify the joint distribution of default times directly, usually through copula structures. For example, a popular static model is the Gaussian copula model proposed by Li (2000), which was the predominant model for the industry to price CDOs (see also Duffie et al., 2003; Andersen et al., 2003). Duffie (2007) points out that the drawbacks of the Gaussian copula model include: (1) The model is internally inconsistent, in that the correlation parameter that matches the price of one tranche of a CDO is typically very different than that of another tranche of the same CDO. This is sometimes called the "correlation smile" (see, e.g. Figure 2 below). (2) The delta hedging of CDO tranches derived from the Gaussian copula model is ineffective. A notorious example of the ineffectiveness occurred with the rate downgrade of General Motors debt in May, 2005, when some market participants who took the delta hedging approach suggested by the Gaussian copula model suffered significant losses. Another major problem with the Gaussian copula model is that it cannot generate tail dependence. More precisely, if the joint distribution of the two default times τ_1 and τ_2 follows a Gaussian copula, it can be shown that

$$\lim_{q \downarrow 0} P(\tau_2 < F_2^{-1}(q) | \tau_1 < F_1^{-1}(q)) = \lim_{q \uparrow 1} P(\tau_2 > F_2^{-1}(q) | \tau_1 > F_1^{-1}(q)) = 0,$$

where F_i is the distribution function of τ_i, $i = 1, 2$ (see Joe, 1997, page 178). For small q, the two events $\{\tau_1 < F_1^{-1}(q)\}$ and $\{\tau_2 < F_2^{-1}(q)\}$ only occur during extreme market conditions, in which strong dependence between the two events is expected, but the Gaussian copula model gives no dependence. For example, conditioning on the rare event of the default of Bear Sterns, the default probability of Lehman Brothers should increase significantly; but in the Gaussian copular model the conditional default probability becomes zero. Therefore, the Gaussian copula model fails to work during financial crises.[1]

The majority of dynamic bottom-up models are intensity based models in which the default time of a single name is modeled as the first jump time of a doubly stochastic Poisson process (Cox process) that is characterized by its default intensity. Single-name intensity based credit risk models are introduced by Jarrow and Turnbull (1995), Lando (1994, 1998), Schönbucher (1998), and Duffie and Singleton (1999). The multi-name intensity based models can be represented by

$$\tau_i = \inf \left\{ t \geq 0 : \int_0^t \lambda_i(s) ds \geq E_i \right\}, i = 1, \ldots, n, \tag{1}$$

where $\lambda_i(t)$ is the default intensity of τ_i; $E_i, i = 1, \ldots, n$ are i.i.d. exponential random variables with mean 1; and $\lambda_i(t)$ is independent of E_j, $\forall i, j = 1, \ldots, n$. Eq. (1) is generally referred to as the *doubly stochastic assumption* in the literature. Under this assumption, the firms' default times τ_i, $i = 1, \ldots, n$ are correlated only as implied by the correlation of their default intensities $\lambda_i(t)$, $i = 1, \ldots, n$, because the default times are conditionally independent, given the filtration generated by the sample paths of default intensities $\lambda_i(t), i = 1, \ldots, n$.

The first multi-name intensity based model for CDO pricing is proposed by Duffie and Gârleanu (2001) and is later extended by Mortensen (2006). More precisely, they postulate that

$$\lambda_i(t) = a_i \lambda^m(t) + \lambda_i^{id}(t), \tag{2}$$

where $\lambda^m(t)$ is the market factor intensity; $\lambda_i^{id}(t)$ is the idiosyncratic intensity of the i-th name; and $\lambda^m(t)$ and $\lambda_i^{id}(t)$, $i = 1, \ldots, n$ are mutually independent. Duffie and Gârleanu (2001)

[1]Extensions to the Gaussian copula model include the double-t copula proposed by Hull and White (2004) and the random factor loading copula proposed by Andersen and Sidenius (2004, 2005) and extended by Andersen (2006), among others.

and Mortensen (2006) model $\lambda^m(t)$ and $\lambda_i^{id}(t)$ as basic affine jump diffusion processes,[2] which provide a significant amount of analytical tractability.[3]

1.2.2 Top-down models

Top-down approaches specify directly the dynamics of the cumulative loss process of the CDO portfolio without modeling the default correlation among individual names. Top-down models are investigated in Arnsdorf and Halperin (2007), Cont and Minca (2007), Errais, Giesecke, and Goldberg (2006), Giesecke and Kim (2007), and Longstaff and Rajan (2007), among others.

1.2.3 Comparison of the Two Categories of Models

Bottom-up models are consistent with underlying single-name default probabilities but generally have more difficulty calibrating to CDO tranche spreads than top-down models. Top-down models can calibrate to CDO tranche spreads but they make little connection to the underlying individual names, because they ignore the relevant information of individual names such as the marginal default probabilities.

In this paper, we propose a dynamic bottom-up model, the conditional survival (CS) model, that can calibrate to both single-name CDS spreads and CDO tranche spreads. Furthermore, the model leads to fast calculation of the sensitivity (the "Greeks") of the CDO tranche spreads with respect to the underlying single-name CDS spreads. Most importantly, the model can produce default clustering that has been empirically observed and is particularly evident during financial crises.

The rest of the paper is organized as follows. The next section discusses the motivation and proposes the CS model. Section 3 provides the details of CDO pricing and the sensitivity analysis of the CDO tranche spreads under the new model. Section 4 shows the numerical results of calibrating the model to the market data on March 14, 2008 and September 16, 2008. Section 5 concludes.

---

[2]The class of affine jump diffusion processes are characterized and analyzed in Duffie, Pan, and Singleton (2000).

[3]Papageorgiou and Sircar (2007) also adopt the framework specified in Eq. (1) and Eq. (2), but instead of using basic affine jump diffusion processes, they use a square-root diffusion process to model $\lambda^m(t)$ and a square-root diffusion process with stochastic volatility to model $\lambda_i^{id}(t)$. Joshi and Stacey (2006) propose the intensity Gamma model $\tau_i = \inf\left\{t \geq 0 : \int_0^t c_i(s)dI(s) \geq E_i\right\}, i = 1, \ldots, n$, where $I(t)$ is a multi-Gamma process representing the market information that drives the default of all names and $c_i(t)$ is a piecewise constant function. Schönbucher (2007) proposes a time-changed intensity model $\tau_i = \inf\left\{t \geq 0 : \int_0^{T(t)} \lambda_i(s)ds \geq E_i\right\}, i = 1, \ldots, n$, where $T(t)$ is a stochastic time change process that introduces correlation among default times.

2 The New Model: Conditional Survival (CS) Model

2.1 Motivation

2.1.1 Default Clustering Effect

The default clustering effect means that one default event tends to trigger more default events both across time and cross-sectionally. The most relevant empirical evidence of the default clustering effect may perhaps come from the recent financial turmoil. Of course, it is still too early to have a systematic empirical investigation of what is happening in the current crisis. However, one can get a feel for the default clustering effect from the recent spread in the CDO market. Table 1 shows the mid bid-ask tranche spreads of the iTraxx Europe Series 8 5-year Index on September 20, 2007 and March 14, 2008, when Federal Reserve provided an emergency loan to near-bankrupt Bear Sterns and financed the purchase of Bear Sterns by JPMorgan Chase two days later. The extremely high spreads of senior tranches on March 14, which are around 10 times as high as they were half a year ago, showed that the default correlation was substantially high and the portfolio loss distribution had significantly heavy tail at the peak of the financial crisis.

Table 1: iTraxx Europe Series 8 5-year Index mid bid-ask tranche spreads on September 20, 2007 and March 14, 2008, quoted in terms of basis points.

Tranches	0%-3%	3%-6%	6%-9%	9%-12%	12%-22%	22%-100%
09/20/2007	1812	84	37	23	15	7
03/14/2008	5150	649	401	255	143	70

There are some empirical evidences of default clustering even before 2007. In terms of default clustering across time, using data on U.S. corporations from 1979 to 2004, Das, Duffie, Kapadia, and Saita (2007) test the joint hypothesis of the doubly stochastic assumption and the well-specified default intensities that are estimated in Duffie, Saita, and Wang (2007). They find some evidence of default clustering exceeding that implied by the doubly stochastic model with the given intensities. In particular, they find that there is a positive serial correlation between the number of defaults in successive time intervals each lasting for some specified amount of time, which supports that defaults cluster in time.

Some traditional stochastic processes, such as Brownian motions and Poisson processes, have independent increments, which makes it hard for them to generate serial correlation. To generate serial correlation, we shall use Pólya process, which is a counting process that is a generalization of a Poisson process. More precisely, a Pólya process $M(t)$ is a Poisson process with a random rate ξ, where ξ is a Gamma random variable with shape parameter α and scale parameter β. In other words, conditional on ξ, $M(t)$ is a Poisson process with rate ξ. The marginal distribution of a Pólya process $M(t)$ is the negative binomial distribution

$$P(M(t) = i) = \binom{\alpha + i - 1}{i} \left(\frac{1}{1 + \beta t}\right)^{\alpha} \left(\frac{\beta t}{1 + \beta t}\right)^{i}, t > 0, i \geq 0,$$

where the binomial coefficient for non-integer $x > y-1, y \geq 0$ is defined by $\binom{x}{y} \triangleq \frac{\Gamma(x+1)}{\Gamma(y+1)\Gamma(x-y+1)}$ with $\Gamma(\cdot)$ being the Gamma function.

A Pólya process has stationary but positively correlated increments. Indeed,

$$\text{Cov}(M(t), M(t + h) - M(t)) = ht\alpha\beta^2 > 0. \tag{3}$$

Hence, the arrival of one event tends to trigger the arrival of more events, which makes Pólya processes suitable for modeling defaults that cluster in time.

In terms of empirical evidence of cross-sectional default clustering, Azizpour and Giesecke (2008) explore the empirical role of contagion, by which the default of a firm has a direct impact on the conditional default rates of the surviving firms, channeled through the complex web of contractual relationships in the economy. For U.S. industrial and financial firms during 1970–2006, they find strong evidence that contagion represents a significant source of default clustering. Longstaff and Rajan (2007) build a three-factor model for the cumulative loss process of the CDX North America Index portfolio, and find that the three-factor model that incorporates the simultaneous defaults of many names has significant incremental explanatory power relative to the one-factor version of the model that incorporates a lower degree of cross-sectional default clustering. Lang and Stulz (1992) provide evidence that a bankruptcy announcement has a contagion effect on other firms in the same industry, by investigating the impact of a firm's bankruptcy announcement upon the stock values of its competitors. Jorion and Zhang (2007) confirm the contagion effect by showing that an extreme upward jump in a firm's CDS spread significantly increase the CDS spreads of its competitors. Collin-Dufresne, Goldstein, and Helwege (2003) find empirically that credit events of large firms generate a market wide increase in credit spreads and downward jumps in risk-free rates.

2.1.2 Motivation for Using Cumulative Default Intensities

A notable fact of default clustering is that many firms can default almost simultaneously, e.g., on the same day. For instance, Azizpour and Giesecke (2008) notice that 1374 corporate defaults were observed in the period from January 1970 to October 2006, while the total number of the days in which at least one default happened was equal to 909; the largest number of events occurred on June 21, 1970, when 24 railway firms defaulted.

However, the model formulated in Eq. (1) is unable to generate simultaneous default of many names. In the model, the default time τ_i is defined as the first passage time of the cumulative default intensity process $\int_0^t \lambda_i(s)ds$ to the random barrier E_i. Even though the default intensity $\lambda_i(t)$ can have jumps, the cumulative default intensity $\int_0^t \lambda_i(s)ds$ is always continuous, because the jump effect is smoothed by the integration. Therefore, even if all the default intensities $\lambda_i(t), i = 1, \ldots, n$ jump together simultaneously, none of the cumulative intensities $\int_0^t \lambda_i(s)ds, i = 1, \ldots, n$ change at all at the time of jump. In face, they only start to increase continuously at higher rates after the jumps occur. Hence, even joint jumps in the default intensities $\lambda_i(t), i = 1, \ldots, n$ cannot generate simultaneous default of many names. This puts a limit on the degree of default clustering that can be incorporated in the model.

One simple way to enhance the capacity of the model to incorporate a high degree of default clustering is to model the cumulative default intensities instead of the instantaneous default intensities, and to allow the cumulative default intensities to have jumps. This motivates the proposal of the new model in the next subsection.

2.2 Conditional Survival (CS) Model

By extending the model of Duffie et al. (2001), we propose the following Conditional Survival (CS) model:

$$\Lambda_i(t) = \sum_{j=1}^{J} a_{i,j} M_j(t) + X_i(t), \tag{4}$$

$$\tau_i = \inf\{t \geq 0 : \Lambda_i(t) \geq E_i\}, i = 1, \ldots, n, \tag{5}$$

where $\Lambda_i(t)$ is the cumulative default intensity of the i-th name. Here $M_j(t)$ represents the j-th market factor in the cumulative default intensities, which is a nonnegative, increasing, and right-continuous stochastic process with $M_j(0) = 0$, $j = 1, \ldots, J$. These market factors can be correlated with each other, but they are all independent of E_i and $X_i(t)$, $\forall i$. Unlike the model formulated in Eq. (1), each $M_j(t)$ is allowed to have jumps. The coefficient $a_{i,j} \geq 0$ is the constant loading of the i-th name on the j-th market factor, $i = 1, \ldots, n; j = 1, \ldots, J$. $X_i(t)$ represents the idiosyncratic part of the cumulative default intensity of the i-th name, which is a nonnegative, right-continuous, and increasing process with $X_i(0) = 0$, $i = 1, \ldots, n$. $X_i(t), i =$

$1, \ldots, n$ are mutually independent, and also independent of the market factors $M_j(t), j = 1, \ldots, J$. $E_i, i = 1, \ldots, n$ are i.i.d. exponential random variables with mean 1; they are all independent of the processes $X_i(t)$ and $M_j(t)$, $\forall i, j$.

Interestingly, there is no need to specify the dynamics of the idiosyncratic default risk factors $X_i(t)$. All that we need to specify in the CS model are the dynamics of the market factors $M_j(t), j = 1, \ldots, J$. We will discuss this more after Proposition 1.

We call the model "conditional survival" because conditional survival probabilities are the building blocks for CDO pricing in the model, as will be shown in Section 3. The conditional survival probabilities in the CS model are very simple. Let $\mathbf{M}(t) \triangleq (M_1(t), \ldots, M_J(t))$ be the vector of market factor processes. Let

$$q_i^c(t) \triangleq P(\tau_i > t | \mathbf{M}(t)) \text{ and } q_i(t) \triangleq P(\tau_i > t) \tag{6}$$

be the conditional survival probability and marginal survival probability of the i-th name, respectively. Then we have the following proposition:

Proposition 1. *For the i-th name in the CS model, we have*

$$q_i^c(t) = E\left[e^{-X_i(t)}\right] e^{-\sum_{j=1}^J a_{i,j} M_j(t)}, \tag{7}$$

$$q_i(t) = E\left[e^{-X_i(t)}\right] E\left[e^{-\sum_{j=1}^J a_{i,j} M_j(t)}\right], \tag{8}$$

And $q_i^c(t)$ can be represented by $q_i(t)$ and market factors as

$$q_i^c(t) = q_i(t) \cdot \frac{e^{-\sum_{j=1}^J a_{i,j} M_j(t)}}{E\left[e^{-\sum_{j=1}^J a_{i,j} M_j(t)}\right]}. \tag{9}$$

Proof. See Appendix A. □

This proposition shows that the conditional survival probability $q_i^c(t)$ can be computed analytically if the marginal survival probability $q_i(t)$ is known and the Laplace transforms of the market factors $\mathbf{M}(t)$ have closed-form formulae. The marginal survival probabilities $q_i(t), i = 1, \ldots, n$ can usually be implied from market quotes of single-name derivatives or from relevant data analysis, since it seems unreasonable to include a name with unknown marginal default probability in the underlying portfolio of a CDO.

The conditional survival probabilities $q_i^c(t), i = 1, \ldots, n$ are useful because given the market factors $\mathbf{M}(t)$, the random variables $1_{\{\tau_i \leq t\}}$, $i = 1, \ldots, n$ are conditionally independent, and $1_{\{\tau_i \leq t\}}$ has a Bernoulli$(1 - q_i^c(t))$ distribution. This enables us to compute CDO tranche spreads.

More precisely, let N_i and R_i be the notional principal and recovery rate of the i-th name in the portfolio, respectively. Then the cumulative loss process of the portfolio is given by

$$L_t = \sum_{i=1}^{n}(1 - R_i)N_i 1_{\{\tau_i \le t\}}, t \ge 0. \tag{10}$$

Therefore, conditional on $\mathbf{M}(t)$, L_t is equal to the linear combination of independent Bernoulli random variables, which further implies that the marginal distribution of L_t is only determined by the dynamics of the market factors $\mathbf{M}(t)$. Since CDO tranche spreads only depend on the marginal distribution of L_t (see Section 3.1), it follows that CDO tranche spreads only depend on the dynamics of the market factors $\mathbf{M}(t)$.

From Eq. (9), the CS model does not need the dynamics of idiosyncratic cumulative intensities $X_i(t)$. Hence no parameters for $X_i(t)$ are needed and there is no need to simulate $X_i(t)$ for pricing CDOs.

2.3 Specifying Market Factors in the Model

According to the previous discussion, to price CDO tranches, we only need to specify the dynamics of the market factors $\mathbf{M}(t)$.[4] We shall use Pólya processes and discrete integral of CIR processes as market factors.

Using a Pólya process as one of the market factors has several advantages:[5] (1) A Pólya process has stationary but positively correlated increments, which makes it suitable for modeling defaults that cluster in time. (2) The jumps of the market factor Pólya process cause simultaneous jumps in the cumulative intensities of all names, which produces strong cross-sectional correlation among individual names. (3) A Pólya process is computationally tractable. The simulation of a Pólya process, which comprises the simulation of a Gamma random variable and a Poisson process, is straightforward. In addition, the Laplace transform of a Pólya process can be calculated in closed form (see Appendix C).

[4] The market factors $\mathbf{M}(t)$ could represent both observable and unobservable macroeconomic variables that have market wide impact on firms' default probability. For example, Duffie et al. (2007) use the trailing one-year return on S&P 500 index and the 3-month Treasury rate to specify firms' default intensities. But using only market observables is not enough. Duffie, Eckner, Horel, and Saita (2008) extend the model in Duffie et al. (2007) by including in the default intensities an unobservable macroeconomic covariate and unobservable firm-specific covariates, and they find strong evidence that firms are exposed to a common dynamic unobservable factor driving default. They show that the unobservable factor induces a much greater probability of extreme portfolio losses than would be the case without the unobservable factor.

[5] A Pólya process is a counting process which has a fixed jump size 1 at each jump time. If one finds that the fixed jump size might be restrictive, one could use a compound Pólya process to incorporate random jump sizes. A compound Pólya process is easy to simulate, and its Laplace transform also has a closed-form formula, which is derived in Appendix C. In the numerical results shown in Section 4, we specify two of the market factors as Pólya processes and find that they seem to be good enough for our purpose.

While we use Pólya processes to produce the high degree of default clustering that typically occurs during financial crises, we shall use a discrete integral of CIR process to model the mild fluctuation of market conditions during normal periods.[6] The discrete integral of CIR process is defined as

$$M(t) \triangleq \frac{h}{2}\lambda(t_0) + h\sum_{i=1}^{m-1}\lambda(t_i) + \frac{h}{2}\lambda(t_m), \tag{11}$$

where m is the number of discretization steps in the time interval $[0,t]$, $h = \frac{t}{m}$, $t_i = \frac{i}{m}t$, $i = 0,\ldots,m$, and $\lambda(t)$ is a CIR process[7] with dynamics

$$d\lambda(t) = \kappa(\theta - \lambda(t))dt + \sigma\sqrt{\lambda(t)}dW(t), \tag{12}$$

where $W(t)$ is a standard Brownian motion. The mean-reverting property of the CIR process makes it suitable to model the fluctuation of macroeconomic conditions. For example, the two macroeconomic covariates proposed in Duffie et al. (2007), i.e., the trailing one-year return on S&P 500 index and the 3-month Treasury rate, are both modeled by mean-reverting time series models. The discrete integral $M(t)$ defined in Eq. (11) is the trapezoidal approximation to the exact integral $\int_0^t \lambda(s)ds$. The advantage of using the discrete integral $M(t)$ is that the exact simulation of $M(t)$ only involves the simulation of the CIR process $\lambda(t)$, which could be easier and faster than the exact simulation of $\int_0^t \lambda(s)ds$ (see Broadie and Kaya, 2006, for the exact simulation of $\int_0^t \lambda(s)ds$). In addition, the Laplace transform of the discrete integral of CIR process has a closed-form formula (see Appendix D).

2.4 Properties of the CS Model

The CS model exhibits some good properties.

[6] Independently in a recent paper, Hull and White (2007) also propose to model the cumulative default intensities, but there are key differences. (1) Their model is a homogeneous one, in which all the names have the same cumulative default intensity $\Lambda(t) = \int_0^t \lambda(s)ds + \sum_{j=1}^{N(t)} H_j$, so it cannot calibrate to individual default probabilities or calculate sensitivities of tranche spreads with respect to individual names. The CS model is a heterogeneous one that can do both. (2) In their model, $\lambda(t)$ is a deterministic function, $N(t)$ is a Poisson process and H_j are deterministic jump sizes. The CS model use Pólya process instead of Poisson process to introduce jumps, and it allows multiple market factors. The CS model does not specify dynamics for idiosyncratic risk factors, but their model use deterministic functions $\int_0^t \lambda(s)ds$ for idiosyncratic risk factor. (3) In Hull and White (2007), it is mentioned that the model can be generalized to the heterogeneous case by allowing different name to have different $\lambda(t)$, and that the heterogeneous case "would be prohibitively time consuming". In contrast, the pricing method for the CS model proposed in Section 3 is tractable enough for calibration. (4) Furthermore, in their heterogeneous case, all the names have the same factor loading 1 on the market factor, but the CS model allows different names to have different loading coefficients on market factors.

[7] The CIR process is an affine diffusion process that is studied by Feller (1951) and is proposed by Cox, Ingersoll, and Ross (1985) as a short rate model, generally referred to as the CIR model. The transition density of the CIR process is known and is related to the noncentral chi-square distribution, so it can be exactly simulated (see Glasserman, 2004).

1. The CS model can generate the simultaneous default of many names, which captures a higher degree of cross-sectional default clustering than can be incorporated in traditional intensity based models as in Eq. (1). More precisely, when a market factor $M_j(t)$ jumps, all cumulative default intensities $\Lambda_i(t), i = 1, \ldots, n$ jump together simultaneously, which can cause many $\Lambda_i(t)$ to cross their respective barriers E_i, i.e., to cause many names to default simultaneously.

2. The CS model can generate defaults that highly cluster in time by incorporating market factors that have jumps clustering in time.

3. The CS model provides automatic calibration to the underlying single-name CDS in the portfolio (see Section 2.5).

4. The CS model does not specify the dynamics of idiosyncratic cumulative intensities $X_i(t)$ in Eq. (4). Hence, no parameters for $X_i(t)$ are needed and there is no need to simulate $X_i(t)$ for pricing CDOs, thanks to Eq. (9).

5. The CS model allows fast CDO pricing based on simulation because one only needs to simulate the value of market factors at coupon payment dates (see Section 3.2).

6. The CS model allows fast calculation of sensitivities (the "Greeks") of CDO tranche spreads with respect to the underlying single-name CDS spreads (see Section 3.4).

The number of free parameters in the CS model is not big (equal to 7 in the numerical examples given in Section 4). The only free parameters in the CS model are the parameters for specifying the dynamics of market factors. The loading coefficients $a_{i,1}, a_{i,2}, \ldots, a_{i,J}$ in Eq. (4) are endogenously determined by a method similar to regression (see Appendix F.1), so they are not free parameters but functions of the market factor parameters.

2.5 Comparison with Other Models

In contrast to other intensity based models, the CS model uses marginal survival probabilities $q_i(t), i = 1, \ldots, n$ as input. This has two advantages: (1) It enables the CS model to automatically calibrate to single-name default probabilities. (2) The CS model does not specify dynamics for idiosyncratic risk factors $X_i(t)$, thanks to Eq. (9) that links the conditional and unconditional survival probabilities without using $X_i(t)$. Therefore the number of parameters in the model is greatly reduced, and fast simulation for CDO pricing becomes feasible because there is no more need to simulate any of the processes $X_i(t)$, $i = 1, \dots, n$. In contrast, $X_i(t)$ is explicitly specified as the integral of basic affine jump diffusion process in Duffie et al. (2001) and Mortensen (2006) and the integral of square-root diffusion with stochastic volatility in Papageorgiou et al. (2007), respectively.

Of course, using marginal survival probabilities $q_i(t), i = 1, \dots, n$ as input to the CS model also incurs some cost: It imposes constraints on the factor loading coefficients $a_{i,j}, j = 1, \dots, J$, which are described in Appendix B. All these constraints are enforced in our numerical examples.

3 CDO Pricing and Sensitivity Analysis in the CS Model

3.1 Synthetic CDO Tranche Spread Formula

Let $T_0 = 0$ be the effective date of the synthetic CDO contract, T be the maturity date of the CDO, and $0 < T_1 < T_2 < \cdots < T_m = T$ be the coupon payment dates. Let L_t be the cumulative loss process defined in Eq. (10) and $[a, b]$ be the CDO tranche loss window. The tranche cumulative loss process is defined as

$$L_t^{[a,b]} = (L_t - a)^+ - (L_t - b)^+, 0 \le t \le T, \tag{13}$$

which denotes the cumulative loss assumed by the investor of the CDO tranche up to time t. Let $D(0, t)$ be the risk free discounter factor from time t to time 0. Then for $a > 0$, the tranche spread $S^{[a,b]}$ is given by

$$S^{[a,b]} = \frac{E\left[\sum_{i=1}^m D(0, \frac{T_{i-1}+T_i}{2})(L_{T_i}^{[a,b]} - L_{T_{i-1}}^{[a,b]})\right]}{E\left[\sum_{i=1}^m D(0, T_i)(T_i - T_{i-1})\left(b - a - \frac{L_{T_i}^{[a,b]}+L_{T_{i-1}}^{[a,b]}}{2}\right)\right]}, \ a > 0. \tag{14}$$

The contractual specification of cash flows for the equity tranche is different from those of the other tranches. The seller of the equity tranche pays an up-front fee at the effective date of the CDO and pays coupons at a fixed running spread of 500 basis points per year to the buyer. The equity tranche spread is defined as the ratio of the up-front fee to the notional of the equity tranche, which is given by

$$\begin{aligned} S^{[0,b]} = \frac{1}{b} \Bigg\{ & E\left[\sum_{i=1}^m D(0, \frac{T_{i-1}+T_i}{2})(L_{T_i}^{[0,b]} - L_{T_{i-1}}^{[0,b]})\right] \\ & -0.05E\left[\sum_{i=1}^m D(0, T_i)(T_i - T_{i-1})\left(b - \frac{L_{T_i}^{[0,b]} + L_{T_{i-1}}^{[0,b]}}{2}\right)\right]\Bigg\}. \end{aligned} \tag{15}$$

The proofs of the above two formulae are given in Appendix E.

3.2 CDO Pricing by Exact Simulation

It is clear from Eq. (14) and (15) that the CDO tranche spreads are determined by the marginal distribution of the cumulative loss process L_t at coupon payment dates T_k, $k = 1, \ldots, m$. Hence, to price CDOs, we only need to simulate L_{T_k} exactly.

An advantage of using marginal survival probabilities as input to the CS model is that it renders the simulation of idiosyncratic risk factors unnecessary. One only needs to simulate the market factors $\mathbf{M}(t)$, the dimension of which is typically very small (equal to 3 in our numerical examples shown in Section 4). Thus, it allows fast simulation for CDO pricing.

The key observation is that conditional on $\mathbf{M}(t)$, the random variables $1_{\{\tau_i \le t\}}$, $i = 1, \ldots, n$ are conditionally independent, and $1_{\{\tau_i \le t\}}$ has a Bernoulli$(1 - q_i^c(t))$ distribution, $i = 1, \ldots, n$.

Suppose $\mathbf{M}(t)$ can be easily simulated and its Laplace transform can be calculated in closed form, then we have the following algorithm to simulate the cumulative loss L_{T_k}, $k = 1, \ldots, m$:

1. Generate sample path of market factors $\mathbf{M}(T_k)$, $k = 1, \ldots, m$.
2. For each $i = 1, \ldots, n$, calculate the conditional survival probability $q_i^c(T_k)$, $k = 1, \ldots, m$, according to Eq. (9), using the samples of market factors generated in Step 1.
3. Generate independent Bernoulli random variables $I_{i,k} \overset{d}{\sim}$ Bernoulli$(1 - q_i^c(T_k))$, $i = 1, \ldots, n$; $k = 1, \ldots, m$.
4. Calculate $L_{T_k} = \sum_{i=1}^{n} (1 - R_i) N_i I_{i,k}$, $k = 1, \ldots, m$.

Using samples of L_{T_k} generated by the above algorithm, we can calculate the estimate and confidence interval for the tranche spreads $S^{[a,b]}$ given in Eq. (14) and (15). By Eq. (10), the expectation of the cumulative loss L_{T_k} is given by $E[L_{T_k}] = \sum_{i=1}^{n} (1 - R_i) N_i [1 - q_i(T_k)]$. So L_{T_k}, $k = 1, \ldots, m$ can be used as the control variates for variance reduction in CDO pricing.

3.3 CDO Pricing without Simulation

We also develop a semi-analytical method for CDO pricing based on the Laplace transform of L_t. The details of the method are available from the authors upon request. The conditional Laplace transform of L_t given $\mathbf{M}(t)$ can be easily calculated because of the conditional independence of the default times. Then the Laplace transform of L_t can be obtained by de-conditioning, i.e., by integrating the conditional Laplace transform with respect to the probability density of $\mathbf{M}(t)$. The Laplace transform of $E[(L_t - a)^+]$ with respect to a, which involves the Laplace transform of L_t, can then be calculated. At last, $E[(L_t - a)^+]$ can be obtained by numerical Laplace inversion of its Laplace transform.

However, in our numerical experiments, we find that the semi-analytical method is not faster than the simulation method because of the need of calculating the density of $\mathbf{M}(t)$ by numerical inversion of its Fourier transform and the need of numerical integration with respect to the density of $\mathbf{M}(t)$. In particular when $\mathbf{M}(t)$ is multi-dimensional, the semi-analytical method loses its advantage.

3.4 Sensitivity Analysis in the CS Model

A prerequisite for active management of the risk of the CDO is the ability to accurately calculate the sensitivities of CDO tranche spreads with respect to market and model parameters, most importantly the single-name CDS spreads of the firms in the underlying portfolio. The number of such sensitivities can be very large, and the calculation of each of these sensitivities can be significantly more challenging than calculating the CDO tranche spreads.

Because the CDS spread of a firm is determined by its marginal survival probabilities, calculating the sensitivities with respect to that single-name CDS spread reduces to calculating the sensitivities with respect to the marginal survival probabilities of that firm. In addition, by Eq. (14) and (15), the CDO tranche spreads are determined by the expected tranche loss $E[L_t^{[a,b]}]$, which is equal to $E[(L_t - a)^+] - E[(L_t - b)^+]$ (cf. Eq. (13)). Therefore, calculating the sensitivities of CDO tranche spreads with respect to single-name CDS spreads reduces to calculating the sensitivities of $E[(L_t - a)^+]$ with respect to marginal survival probabilities $q_i(t)$.

One of the major advantages of the CS model is that the sensitivities of CDO tranche spreads with respect to each of the n single-name CDS spreads can be obtained at the same time when the tranche spreads are calculated by simulation, with little extra computation effort. More precisely, we have the following proposition:

Proposition 2. *The sensitivities of $E[(L_t - a)^+]$ with respect to the survival probability of the i-th name is*

$$\frac{\partial E[(L_t - a)^+]}{\partial q_i(t)} =$$
$$E\left\{ \frac{e^{-\sum_{j=1}^{J} a_{i,j} M_j(t)}}{E\left[e^{-\sum_{j=1}^{J} a_{i,j} M_j(t)}\right]} \left[(L_t^{(-i)} - a)^+ - (L_t^{(-i)} + (1 - R_i)N_i - a)^+\right] \right\}, \tag{16}$$

where $L_t^{(-i)} \triangleq \sum_{j\neq i}(1-R_j)N_j 1_{\{\tau_j \leq t\}}$ *represents the cumulative loss of the portfolio excluding the i-th name. In particular, at time* $t = T_k$, $L_{T_k}^{(-i)} = \sum_{j\neq i}(1-R_j)N_j I_{j,k}$, *where* $I_{j,k}, j = 1,\ldots,n$ *are the Bernoulli random variables generated in the pricing algorithm in Section 3.2.*

Proof. See Appendix A. □

In view of this proposition, the sensitivities can be calculated at the same time as the CDO tranches are priced. Of course, one can also use $L_{T_k}, k = 1,\ldots,m$ as control variates for variance reduction in calculating the sensitivities.

4 Numerical Results

In this section, we show the numerical results of calibrating the CS model to both the iTraxx Europe 5-year CDO Index tranche spreads and the individual CDS spreads of the underlying firms on two dates, March 14, 2008 (right before the collapse of Bear Sterns) and September 16, 2008 (right after Lehman Brothers went bankrupt). All the market data are obtained from Bloomberg business electronic resources.

4.1 Description of the Data

The iTraxx Europe index consists of the most liquid 125 CDS referencing European investment-grade firms. Every six months in March and September a new series of index is launched and the firms included in the index portfolio are updated, with some illiquid or downgraded firms being dropped and new ones being added. The most recently launched series of index is called "on-the-run" index. We shall use the on-the-run index, as it is more liquid than previously issued indices, which are known as off-the-run indices; this is especially true during the current crisis in which there are very few trades for the off-the-run indices.

The iTraxx Europe index has 6 tranches that correspond to the 0%-3%, 3%-6%, 6%-9%, 9%-12%, 12%-22%, and 22%-100% of the portfolio notional, respectively. The market spread of the 0%-3% tranche, i.e., the equity tranche, is quoted in terms of percentage up-front fee, while those of other tranches are quoted in terms of running spreads. All the tranches pay quarterly coupons.

The iTraxx Europe Series 8 5-year Index was launched on September 20, 2007 and has a maturity of 5 years. It was the on-the-run index on March 14, 2008. The underlying portfolio of the index consists of 124 firms. Each firm has a notional amount of Euro 8/3 million except UniCredit SpA which has a notional of Euro 16/3 million. The first coupon payment date after

March 14, 2008 is June 20, 2008. The last coupon payment date is December 20, 2012. The iTraxx Europe Series 9 5-year Index, which was launched on March 20, 2008, was the on-the-run index on September 16, 2008. The underlying portfolio of the Series 9 index also comprises 124 firms, 112 of which were included in the underlying portfolio of the Series 8 index. Each firm in the portfolio has a notional of Euro 8/3 million except GDF SUEZ which has a notional of Euro 16/3 million. The first coupon payment date after September 16, 2008 is December 22, 2008. The last coupon payment date is June 20, 2013.

To fully calibrate the model to individual CDS spreads, we also use the CDS data of the 124 firms in the on-the-run iTraxx Europe 5-year Index portfolio on March 14, 2008 and September 16, 2008. Table 2 shows the summary statistics of the mid bid-ask 5-year CDS spreads of the firms in the two portfolios.

Table 2: Summary statistics of the closing data of the mid bid-ask 5-year CDS spreads of the 124 firms in the iTraxx Europe Series 8 5-year Index portfolio on March 14, 2008, and those of the 124 firms in the iTraxx Europe Series 9 5-year Index portfolio on September 16, 2008, in terms of basis points.

Statistics	03/14/2008	09/16/2008
Min	61.50	27.20
Max	489.70	495.80
Mean	163.57	140.09
Median	150.10	117.00
Standard deviation	76.98	79.88

Table 3 shows the discount factors that discount from the coupon payment dates of iTraxx Europe 5-year index tranches back to March 14, 2008 and September 16, 2008. These discount factors are extracted from the Euro fixing swap curve on those two dates, and they will be used in the tranche spread formulae (14) and (15).

4.2 Objective Function of Calibration

Let $s_k^{o,b}$ and $s_k^{o,a}$ be the bid and ask of the k-th CDO tranche spread observed in the market, respectively, $k = 1, \ldots, K$, where K is the number of tranches. Let $s_k^o = (s_k^{o,a} + s_k^{o,b})/2$ be the mid bid-ask tranche spread. Let s_k be the tranche spread computed by the CS model. The chi-square test statistic is given by

$$\chi^2 = \sum_{k=1}^{K} \frac{(s_k - s_k^o)^2}{s_k}. \tag{17}$$

Table 3: Discount factors on March 14, 2008 and September 16, 2008

03/14/2008		09/16/2008	
Coupon Date	Discount Factor	Coupon Date	Discount Factor
03/14/2008	1.0000	09/16/2008	1.0000
06/20/2008	0.9878	12/22/2008	0.9868
09/22/2008	0.9762	03/20/2009	0.9741
12/22/2008	0.9653	06/22/2009	0.9609
03/20/2009	0.9551	09/21/2009	0.9481
06/22/2009	0.9471	12/21/2009	0.9392
09/21/2009	0.9403	03/22/2010	0.9314
12/21/2009	0.9319	06/21/2010	0.9217
03/22/2010	0.9238	09/20/2010	0.9123
06/21/2010	0.9152	12/20/2010	0.9023
09/20/2010	0.9068	03/21/2011	0.8926
12/20/2010	0.8985	06/20/2011	0.8830
03/21/2011	0.8903	09/20/2011	0.8735
06/20/2011	0.8817	12/20/2011	0.8638
09/20/2011	0.8731	03/20/2012	0.8543
12/20/2011	0.8646	06/20/2012	0.8448
03/20/2012	0.8562	09/20/2012	0.8355
06/20/2012	0.8474	12/20/2012	0.8260
09/20/2012	0.8387	03/20/2013	0.8167
12/20/2012	0.8301	06/20/2013	0.8073

The root-mean-square error of the model is given by

$$\mathrm{RMSE} = \sqrt{\frac{1}{K}\sum_{k=1}^{K}\left(\frac{s_k - s_k^o}{s_k^{o,a} - s_k^{o,b}}\right)^2}. \tag{18}$$

We use the chi-square test statistic as the objective function in the calibration, i.e., we search for model parameters that minimize the chi-square test statistic. We also report the root-mean-square error as a reference. In the numerical examples, we find that there is not much difference between using chi-square test statistic and using root-mean-square error as the objective function.

4.3 Calibration Results

We use three independent market factors in the CS model for calibration: two Pólya processes and one discrete integral of CIR process. For any $c > 0$ and a CIR process $\lambda(t)$ with parameters $(\kappa, \theta, \sigma, \lambda(0))$, the dynamics of which is given by Eq. (12), we note that $c\lambda(t)$ is still a CIR process with parameters $(\kappa, c\theta, \sqrt{c}\sigma, c\lambda(0))$. Therefore, in the CS model formulated in Eq. (4) and (5), one can always make the long-run mean of the CIR process equal to a pre-defined constant θ_0 without changing the model, by multiplying the CIR process by a factor θ_0/θ and dividing the corresponding loading coefficient $a_{i,j}$ by θ_0/θ at the same time. In other words, only three free parameters $(\kappa, \sigma, \lambda(0))$ need to be determined for the CIR process. In the calibration, we choose $\theta_0 = 0.1$. Hence, there are 7 market factor parameters $\Theta = (\alpha_1, \beta_1, \alpha_2, \beta_2, \kappa, \sigma, \lambda(0))$ in total, where α_i and β_i are parameters for the i-th Pólya process, $i = 1, 2$, and $\kappa, \sigma, \lambda(0)$ are the parameters for the CIR process. For the discrete integral of CIR processes, we use 10 discretization steps for the first coupon payment period (e.g., from March 14, 2008 to June 20, 2008 for iTraxx Series 8 data on March 14, 2008), and 8 steps for each of the remaining coupon payment periods.

We use the market quotes of the mid bid-ask 5 year CDS spreads of the 124 firms to extract the marginal survival probabilities $q_i(t), i = 1, \ldots, 124$ for each firm, assuming that the hazard rate functions are flat during the 5 year period. We assume fixed recovery rate $R_i = 40\%$, $i = 1, \ldots, n$. The details of the calibration procedure are described in Appendix F.

We carry out a joint calibration of the data on March 14, 2008 and the data on September 16, 2008. We would like the two sets of parameters that correspond to the two sets of data to be as close to each other as possible, so that we can find out which of the 7 parameters in the CS model tend to be stable with respect to time. In the joint calibration, we keep the 3 CIR parameters $(\kappa, \sigma, \lambda(0))$ for the data on March 14, 2008 and those for the data on September 16, 2008 to be the same.

Table 4 shows the calibrated parameters and the summary of the calibration results. Both p-values corresponding to the chi-square test statistics for the two sets of data are larger than 5%, indicating a reasonable fit of the CS model, even though the data correspond to two dates that are half a year apart, and the portfolios underlying the two indices are not exactly the same. This shows the stability of the parameters related to the CIR factor. The two parameters α_1 and β_2 for the Pólya processes do not change much, too. The other two parameters β_1 and α_2 seem to vary significantly in terms of their respective relative changes.

By Eq. (3), the positive correlation between the increments of a Pólya process with parameters α and β is determined by $\alpha\beta^2$. Therefore, it is clear from Table 4 that the first Pólya process (with parameters α_1 and β_1) has much weaker correlation between its increments than the second Pólya process (with parameters α_2 and β_2).

Table 5 shows the details of the calibration results for the two sets of data using parameters shown in Table 4.

The iTraxx Europe Series 9 Index data on September 16, 2008 provides an example that the Gaussian copula model does not work during financial crises, as the implied copula correlation for the 6%-9% tranche does not exist! In other words the market quote of tranche spread

Table 4: Calibrated parameters for the on-the-run iTraxx Europe 5-year Index tranche spreads on March 14, 2008 and September 16, 2008. The CIR parameters $(\kappa, \sigma, \lambda(0))$ are the same for the two sets of data in the calibration. The last three rows show the χ^2 test statistic (defined in Eq. (17)), the corresponding p-value, and the root-mean-square error (defined in Eq. (18)) of the calibration.

	03/14/2008	09/16/2008
α_1	0.53308071875339	0.68644645282521
β_1	0.03037513018266	0.01800593339554
α_2	0.00215211534704	0.00578696362877
β_2	8.28982581200579	9.02448448266147
κ	0.05260528397600	
σ	1.68370042003610	
$\lambda(0)$	1.91761310257449	
χ^2	8.71	9.80
p-value	0.12	0.08
RMSE	1.24	2.66

Table 5: Calibration result for on-the-run iTraxx Europe 5-year Index tranche spreads on March 14, 2008 and September 16, 2008, using parameters shown in Table 4. We use 50,000 replications in the Monte Carlo (MC) simulation. All the spreads are expressed in terms of basis points. The χ^2 test statistic is defined in Eq. (17). The root-mean-square error is defined in Eq. (18).

March 14, 2008						
Tranches	0%-3%	3%-6%	6%-9%	9%-12%	12%-22%	22%-100%
Market spreads	5149.95	649.00	401.13	255.31	143.40	69.90
Model spreads	5048.95	691.14	395.47	261.23	168.62	66.96
Bid-ask spreads	158.10	24.38	24.59	19.79	11.75	2.92
MC standard error	15.38	5.07	3.84	3.18	2.60	1.13
$\chi^2 = 8.71$ (p-value = 0.12), RMSE = 1.24						
September 16, 2008						
Tranches	0%-3%	3%-6%	6%-9%	9%-12%	12%-22%	22%-100%
Market spreads	4598.00	618.25	374.50	215.16	102.17	58.81
Model spreads	4610.30	630.49	347.43	217.43	131.53	52.18
Bid-ask spreads	118.00	14.00	12.50	10.55	5.33	2.58
MC standard error	15.85	4.81	3.55	2.87	2.29	1.00
$\chi^2 = 9.80$ (p-value = 0.08), RMSE = 2.66						

(375 bp) is not in the range of spreads that can be produced by the Gaussian copula model. In fact, when the correlation parameter in the Gaussian copula model increases from 0% to 100%, the corresponding 6%-9% tranche spread generated by the copula model first increases from 488 bp to 719 bp, then decreases from 719 bp to 400 bp. To keep the integrality of the figure, we use 100% as a substitute for the implied correlation of the 6%-9% tranche, since 400 bp (corresponding to the correlation parameter 100%) is the closest to the observed tranche spread (375 bp).

Fig. 2 shows the implied Gaussian copula correlation of the market quotes on March 14, 2008 and September 16, 2008, and the implied correlation of the tranche spreads generated by the CS model, using parameters shown in Table 4. The skew of the implied copula correlation, which is dramatic during the financial crisis, is excellently reproduced by the CS model.

5 Conclusion

The current subprime mortgage crisis has demonstrated the importance of studying models that can incorporate the default clustering effect into the pricing of CDOs. In this paper, we propose a conditional survival (CS) model that can generate a substantially high degree of default clustering. By using marginal survival probabilities as input, the CS model automatically calibrates to single-name CDS spreads, and the model allows fast CDO pricing and calculation of sensitivities of CDO tranche spreads with respect to the underlying single-name CDS spreads. This provides a direct link between single-name CDS and multi-name CDOs. The result of calibration to the recent market data on March 14, 2008 and September 16, 2008, when major financial institutions collapsed and default correlation among firms was substantially high, shows that the model is promising.

A Proof of Propositions

Proof of Proposition 1. By Eq. (5), we have

$$
\begin{aligned}
q_i^c(t) &= E\left[1_{\{X_i(t)+\sum_{j=1}^J a_{i,j}M_j(t)<E_i\}}\middle|\mathbf{M}(t)\right] \\
&= E\left[E\left[1_{\{X_i(t)+\sum_{j=1}^J a_{i,j}M_j(t)<E_i\}}\middle|X_i(t),\mathbf{M}(t)\right]\middle|\mathbf{M}(t)\right] \\
&= E\left[e^{-X_i(t)-\sum_{j=1}^J a_{i,j}M_j(t)}\middle|\mathbf{M}(t)\right] \\
&= E\left[e^{-X_i(t)}\middle|\mathbf{M}(t)\right]e^{-\sum_{j=1}^J a_{i,j}M_j(t)}
\end{aligned}
$$

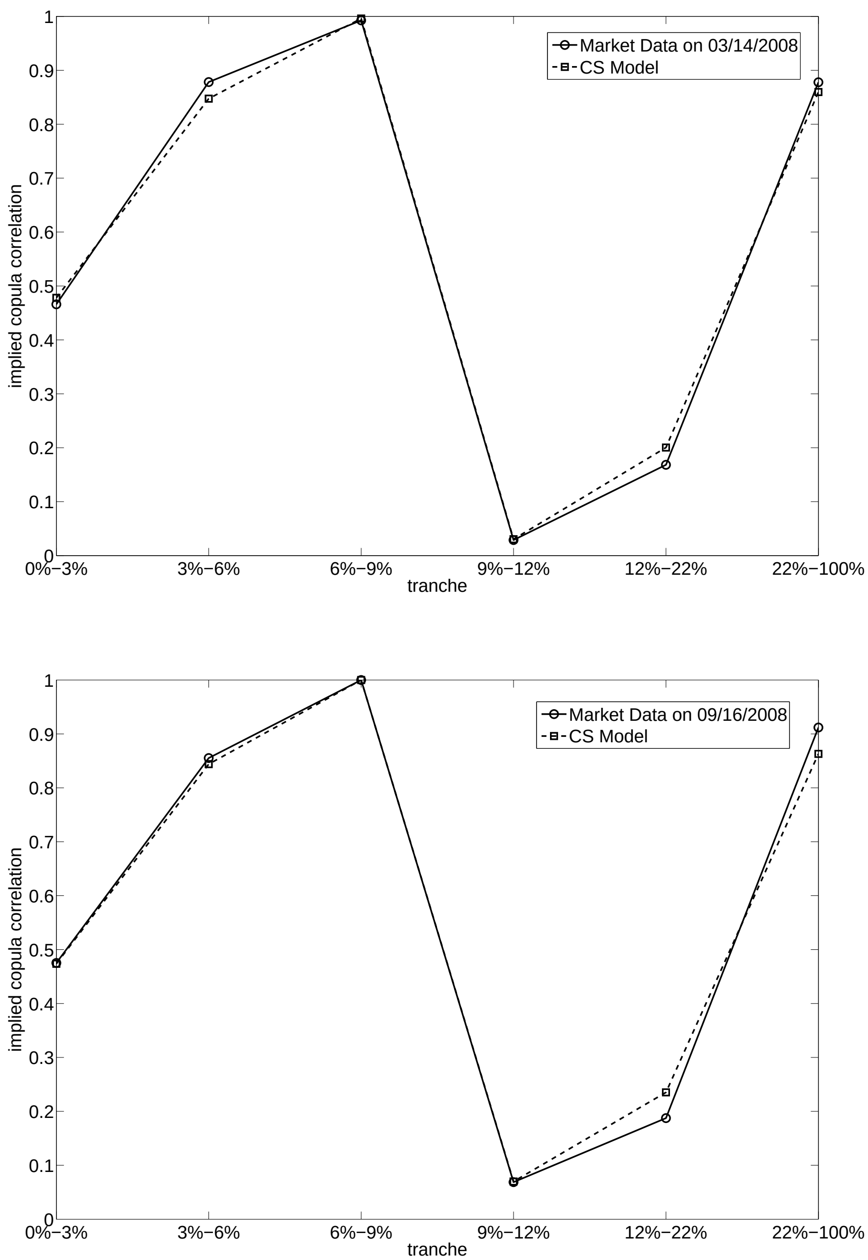

Figure 2: Comparison between the implied copula correlation of the CS model and that of the market quotes of the on-the-run iTraxx Europe 5-year Index on March 14, 2008 (upper panel) and September 16, 2008 (lower panel). Note that on September 16, 2008 the market implied copula correlation of the 6%-9% tranche does not exist; so we use 100% as a substitute.

$$=E\left[e^{-X_i(t)}\right]e^{-\sum_{j=1}^{J}a_{i,j}M_j(t)}, \tag{19}$$

where the last equality follows from the independence of $X_i(t)$ and $\mathbf{M}(t)$. Then by taking expectation on both sides of Eq. (19), we obtain Eq. (8). Dividing Eq. (7) by (8), we obtain Eq. (9). □

Proof of Proposition 2. Let $c_i \triangleq (1-R_i)N_i$. Noting that conditional on $\mathbf{M}(t)$, the n default indicators $1_{\{\tau_i\le t\}}$, $i=1,\dots,n$ are conditionally independent, we have

$$\begin{aligned}
&E[(L_t-a)^+|\mathbf{M}(t)] = E[(L_t^{(-i)}+c_i1_{\{\tau_i\le t\}}-a)^+|\mathbf{M}(t)]\\
=&E[E[(L_t^{(-i)}+c_i1_{\{\tau_i\le t\}}-a)^+|\mathbf{M}(t),L_t^{(-i)}]|\mathbf{M}(t)]\\
=&E[(L_t^{(-i)}+c_i-a)^+(1-q_i^c(t))+(L_t^{(-i)}-a)^+q_i^c(t)|\mathbf{M}(t)]\\
=&E[(L_t^{(-i)}+c_i-a)^+|\mathbf{M}(t)]\\
&+E\{q_i^c(t)[(L_t^{(-i)}-a)^+-(L_t^{(-i)}+c_i-a)^+]|\mathbf{M}(t)\}.
\end{aligned}$$

Therefore,

$$\begin{aligned}
&E[(L_t-a)^+] = E[(L_t^{(-i)}+c_i-a)^+]+\\
&\quad q_i(t)E\left\{\frac{q_i^c(t)}{q_i(t)}[(L_t^{(-i)}-a)^+-(L_t^{(-i)}+c_i-a)^+]\right\}.
\end{aligned} \tag{20}$$

By Eq. (9), $\frac{q_i^c(t)}{q_i(t)} = \frac{e^{-\sum_{j=1}^{J}a_{i,j}M_j(t)}}{E\left[e^{-\sum_{j=1}^{J}a_{i,j}M_j(t)}\right]}$ does not depend on $q_i(t)$. Thus Eq. (16) is obtained through differentiating the above equation by $q_i(t)$. □

B Constraints on the Loading Coefficients in the CS Model

In the CS model, the idiosyncratic cumulative default intensity $X_i(t)$ is nonnegative and increasing. Therefore, for a sequence of coupon payment dates $0=T_0<T_1<\cdots<T_m=T$, it must hold that

$$E\left[e^{-X_i(T_k)}\right]\le E\left[e^{-X_i(T_{k-1})}\right], 1\le k\le m. \tag{21}$$

By Eq. (8) and (21), we obtain the constraints on the loading coefficients $a_{i,j}$: for each i, we must have

$$\frac{q_i(T_k)}{E\left[e^{-\sum_{j=1}^{J}a_{i,j}M_j(T_k)}\right]} \le \frac{q_i(T_{k-1})}{E\left[e^{-\sum_{j=1}^{J}a_{i,j}M_j(T_{k-1})}\right]}, 1\le k\le m, \tag{22}$$

$$0 \le a_{i,j},\ j=1,\dots,J. \tag{23}$$

C Laplace Transform of a Compound Pólya Process

Let $P(t)$ be a Pólya process with parameters α and β. Then the Laplace transform of $P(t)$ is given by

$$E\left[e^{-uP(t)}\right] = \left(\frac{p}{1-(1-p)e^{-u}}\right)^{\alpha},\ p = \frac{1}{1+\beta t}. \tag{24}$$

Let Y_1, Y_2, ... be i.i.d. random variables with Laplace transform $\phi_Y(u) = E[e^{-uY_1}]$ that are independent of $P(t)$. Then the Laplace transform of the compound Pólya process $M(t) \triangleq \sum_{i=1}^{P(t)} Y_i$ is given by

$$E\left[e^{-uM(t)}\right] = \left(\frac{p}{1-(1-p)\phi_Y(u)}\right)^{\alpha},\ p = \frac{1}{1+\beta t}. \tag{25}$$

Indeed, we can show this by noting that

$$\begin{aligned}
&E\left[e^{-uM(t)}\right] = E\left[E\left[e^{-uM(t)}\middle| P(t)\right]\right] \\
=&E\left[\left(E\left[e^{-uY_1}\right]\right)^{P(t)}\right] = E\left[e^{-\log\left(\frac{1}{\phi_Y(u)}\right)P(t)}\right] \\
=& E\left[e^{-vP(t)}\right]\Big|_{v=\log\left(\frac{1}{\phi_Y(u)}\right)} \\
=&\left(\frac{p}{1-(1-p)\phi_Y(u)}\right)^{\alpha},\ p = \frac{1}{1+\beta t},
\end{aligned}$$

where the last equality follows from Eq. (24).

D Laplace Transform of Discrete Integral of CIR Process

The transition law of $\lambda(t)$ given $\lambda(s)$ for $s < t$ is given by (see Cox et al., 1985) $\lambda(t)|\lambda(s) \overset{d}{\sim} \alpha(s,t)\chi_d'^2(\beta(s,t)\lambda(s))$, where $\chi_v'^2(\lambda)$ denotes the noncentral chi-square distribution with v degrees of freedom and noncentrality parameter λ,

$$\alpha(s,t) = \frac{\sigma^2(1-e^{-\kappa(t-s)})}{4\kappa},\ \ \beta(s,t) = \frac{4\kappa e^{-\kappa(t-s)}}{\sigma^2(1-e^{-\kappa(t-s)})},\ \ d = \frac{4\kappa\theta}{\sigma^2}.$$

The transition law of the CIR process implies that the conditional Laplace transform of $\lambda(t)$ is given by

$$E\left[e^{-u\lambda(t)}\middle| \lambda(s)\right] = (1+2\alpha(s,t)u)^{-\frac{d}{2}} e^{-\frac{\alpha(s,t)\beta(s,t)u}{1+2\alpha(s,t)u}\lambda(s)}. \tag{26}$$

Let $\alpha_i = \alpha(t_i, t_{i+1})$, $\beta_i = \beta(t_i, t_{i+1})$, $i = 0, 1, \ldots, m-1$. Then the Laplace transform of the discrete integral of CIR process defined in (11) is given by

$$E\left[e^{-uM(t)}\right] = e^{f_0(u)\lambda(t_0)}\left(\prod_{i=0}^{m-1}(1-2\alpha_i f_{i+1}(u))\right)^{-d/2}, \tag{27}$$

where $d = \frac{4\kappa\theta}{\sigma^2}$ and $f_m(u)$, $f_{m-1}(u)$, ..., $f_0(u)$ are defined by the recursion

$$\begin{aligned}
f_m(u) &= -\frac{h}{2}u; \\
f_{i-1}(u) &= -hu + \frac{\alpha_{i-1}\beta_{i-1}f_i(u)}{1-2\alpha_{i-1}f_i(u)}, m \geq i \geq 2; \\
f_0(u) &= -\frac{h}{2}u + \frac{\alpha_0\beta_0 f_1(u)}{1-2\alpha_0 f_1(u)}.
\end{aligned} \tag{28}$$

To prove Eq. (27) and (28), let $\mathcal{F}_t \triangleq \sigma(\lambda(s), s \leq t)$ and $f_k(u)$ be defined in Eq. (28), $k = m, \ldots, 0$. Then

$$\begin{aligned}
E\left[e^{-uM(t)}\right] &= E\left[e^{-u(\frac{h}{2}\lambda(t_0)+h\sum_{i=1}^{m-1}\lambda(t_i)+\frac{h}{2}\lambda(t_m))}\right] \\
=&E\left[e^{-u(\frac{h}{2}\lambda(t_0)+h\sum_{i=1}^{m-1}\lambda(t_i))+f_m(u)\lambda(t_m)}\right] \\
=&E\left[e^{-u(\frac{h}{2}\lambda(t_0)+h\sum_{i=1}^{m-1}\lambda(t_i))}\ E\left[e^{f_m(u)\lambda(t_m)}\middle|\ \mathcal{F}_{t_{m-1}}\right]\right] \\
=&E\left[e^{-u(\frac{h}{2}\lambda(t_0)+h\sum_{i=1}^{m-1}\lambda(t_i))}\ E\left[e^{f_m(u)\lambda(t_m)}\middle|\ \lambda(t_{m-1})\right]\right] \\
=&E\left[e^{-u(\frac{h}{2}\lambda(t_0)+h\sum_{i=1}^{m-1}\lambda(t_i))}\frac{e^{\frac{\alpha_{m-1}\beta_{m-1}f_m(u)}{1-2\alpha_{m-1}f_m(u)}\lambda(t_{m-1})}}{(1-2\alpha_{m-1}f_m(u))^{\frac{d}{2}}}\right] \quad \text{(by Eq. (26))} \\
=&(1-2\alpha_{m-1}f_m(u))^{-\frac{d}{2}}E\left[e^{-u(\frac{h}{2}\lambda(t_0)+h\sum_{i=1}^{m-2}\lambda(t_i))+f_{m-1}(u)\lambda(t_{m-1})}\right].
\end{aligned}$$

Repeating the above argument by successively conditioning on $\mathcal{F}_{t_{m-2}}$, $\mathcal{F}_{t_{m-3}}$, ..., and $\mathcal{F}_{t_1}$, we obtain $E\left[e^{-uM(t)}\right] = e^{f_0(u)\lambda(t_0)}\left(\prod_{i=0}^{m-1}(1-2\alpha_i f_{i+1}(u))\right)^{-d/2}$.

E Proving CDO Tranche Spreads Formulae

If a default event happens at time τ, the investor of the tranche would make a payment equal to $L_\tau^{[a,b]} - L_{\tau-}^{[a,b]}$, where $L_{\tau-}^{[a,b]} \triangleq \lim_{t\uparrow\tau} L_t^{[a,b]}$. Therefore, the present value of the default leg of the tranche is equal to $E[\int_0^T D(0,t)dL_t^{[a,b]}]$, under the risk-neutral probability measure. For simplicity, it is usually assumed in the literature that defaults only occur in the middle of coupon payment dates (see e.g. Mortensen, 2006; Andersen et al., 2003; Papageorgiou et al., 2007). Under this simplification, the present value of the default leg is given by

$$\begin{aligned}
E\left[\int_0^T D(0,t)dL_t^{[a,b]}\right] &= E\left[\sum_{i=1}^m \int_{T_{i-1}}^{T_i} D(0,t)dL_t^{[a,b]}\right] \\
&= E\left[\sum_{i=1}^m D(0,\frac{T_{i-1}+T_i}{2})(L_{T_i}^{[a,b]} - L_{T_{i-1}}^{[a,b]})\right].
\end{aligned} \tag{29}$$

Next, we compute the present value of the premium leg. The outstanding notional of the tranche at time t is $O_t^{[a,b]} = b - a - L_t^{[a,b]}$. Then the coupon payment at time $T_i, i = 1, \ldots, m$ is specified as

$$S^{[a,b]} \cdot (T_i - T_{i-1}) \int_{T_{i-1}}^{T_i} \frac{O_t^{[a,b]}}{T_i - T_{i-1}} dt = S^{[a,b]} \int_{T_{i-1}}^{T_i} O_t^{[a,b]} dt.$$

Assuming defaults only occur in the middle of premium periods, the present value of the premium leg is given by

$$\begin{aligned}
& E\left[\sum_{i=1}^{m} D(0,T_i) S^{[a,b]} \int_{T_{i-1}}^{T_i} O_t^{[a,b]} dt\right] \\
= \; & E\left[\sum_{i=1}^{m} D(0,T_i) S^{[a,b]} (T_i - T_{i-1}) \frac{1}{2} (O_{T_{i-1}}^{[a,b]} + O_{T_i}^{[a,b]})\right] \\
= \; & S^{[a,b]} \cdot E\left[\sum_{i=1}^{m} D(0,T_i)(T_i - T_{i-1}) \left(b - a - \frac{L_{T_i}^{[a,b]} + L_{T_{i-1}}^{[a,b]}}{2}\right)\right]. \qquad (30)
\end{aligned}$$

By making the present value of the premium leg and that of the default leg equal, we obtain the fair spread of the tranche

$$S^{[a,b]} = \frac{E\left[\sum_{i=1}^{m} D(0, \frac{T_{i-1}+T_i}{2})(L_{T_i}^{[a,b]} - L_{T_{i-1}}^{[a,b]})\right]}{E\left[\sum_{i=1}^{m} D(0,T_i)(T_i - T_{i-1}) \left(b - a - \frac{L_{T_i}^{[a,b]} + L_{T_{i-1}}^{[a,b]}}{2}\right)\right]}, \text{for } a > 0.$$

The formula for the spread of equity tranche follows as $b \cdot S^{[0,b]}$ is the difference between two terms, the first being the present value of the default leg, and the second being the present value of the premium leg with 500 basis point spread.

F Calibrate the Model to Market CDS and CDO Spreads

Since we use $q_i(t), i = 1, \ldots, n$ implied from market CDS spreads of the firms in the CDO portfolio as input to the CS model, it automatically calibrates to the single-name CDS spreads.

To calibrate to market CDO tranche spreads, we need to determine two sets of parameters: (1) the parameters that specify the dynamics of $\mathbf{M}(t)$, which we denote by Θ; (2) the factor loading coefficients $(a_{i,1}, \ldots, a_{i,J})$ for each name i, $i = 1, \ldots, n$.

F.1 Determine Factor Loading Coefficients by Regression

Suppose the parameters Θ for market factors have been fixed. By Eq. (4), the factor loading coefficients $(a_{i,1}, \ldots, a_{i,J})$ can be viewed as the regression coefficients of $\Lambda_i(t)$ on the regressors $(M_1(t), \ldots, M_J(t))$, with $X_i(t)$ being the regression error term. When the regression error term $X_i(t)$ is small, $E\left[e^{-X_i(t)}\right]$ is close to 1. Therefore, the regression error can be measured by $\left|1 - E\left[e^{-X_i(t)}\right]\right|$. Expecting that good market factors are able to explain a significant part of the default risk, we determine $(a_{i,1}, \ldots, a_{i,J})$ by minimizing the regression error $\left|1 - E\left[e^{-X_i(t)}\right]\right|$. By Eq. (8) and the fact that $E\left[e^{-X_i(t)}\right] \leq 1$ since $X_i(t) \geq 0$, minimizing $\left|1 - E\left[e^{-X_i(t)}\right]\right|$ is equivalent to minimizing $E\left[e^{-\sum_{j=1}^{J} a_{i,j} M_j(t)}\right] - q_i(t)$.

Recalling that the loading coefficients must satisfy the constraints in (22) and (23), we determine $(a_{i,1}, \ldots, a_{i,J})$ by solving the following optimization problem:

$$\begin{aligned} \min \ & E\left[e^{-\sum_{j=1}^{J} a_{i,j} M_j(T)}\right] - q_i(T) \\ \text{s.t.} \ & \frac{q_i(T_k)}{E\left[e^{-\sum_{j=1}^{J} a_{i,j} M_j(T_k)}\right]} \leq \frac{q_i(T_{k-1})}{E\left[e^{-\sum_{j=1}^{J} a_{i,j} M_j(T_{k-1})}\right]}, 1 \leq k \leq m \\ & 0 \leq a_{i,j}, \ j = 1, \ldots, J. \end{aligned} \tag{31}$$

The idea of minimizing the contribution of idiosyncratic risk factors to firms' default probabilities is not inconsistent with the existing studies that find the importance of firm-specific covariates on a firm's default probability, because the "idiosyncratic" risk factors in the CS model are different from the "firm-specific" risk factors commonly referred to in the literature. The former are by definition independent of market factors, but the latter can have significant correlation with market factors. For example, the two firm-specific covariates used in Duffie et al. (2007), i.e., the distance-to-default and the trailing one-year stock return, are both correlated with macroeconomic market factors. Indeed, the correlation is explicitly captured in the time series model in Duffie et al. (2007). Therefore, minimizing the effect of idiosyncratic risk factors does not amount to minimizing that of firm-specific factors.

F.2 The Calibration Algorithm

Using the chi-square test statistic defined in Eq. (17) as the objective function, the calibration algorithm for market factor parameters Θ is as follows:

1. We start from an initial guess of the market factor parameters Θ.

2. Given the market factor parameters Θ, for each $i = 1, 2, \ldots, n$, we determine the factor loading coefficients $(a_{i,1}, \ldots, a_{i,J})$ by solving the problem formulated in (31).

3. Calculate CDO tranche spread $s_k, k = 1, \ldots, K$ in the CS model using the market factor parameters Θ and the factor loading coefficients obtained in the previous step.

4. Calculate the chi-square statistic that corresponds to the current market factor parameters Θ.

5. If the chi-square statistic is small enough, stop; otherwise, update the market factor parameters Θ using some unconstrained optimization algorithm (e.g., Powell's direction-set method, see Press et al. 2002), and go to Step 2.

In step 2, one needs to solve n optimization problems formulated in (31) to find the loading coefficients for all the n names. The dimension of the optimization problem is equal to that of the vector $\mathbf{M}(t)$, which is typically not greater than 3. The objective function and the constraints all have closed-form derivatives, and the objective function is monotonic. Therefore, the optimization problem can be solved quickly by some common methods such as sequential quadratic programming. We used the CFSQP package developed by Lawrence et al. (1997) in our implementation.

References

Andersen, L., 2006. Portfolio losses in factor models: term structures and intertemporal loss dependence. Working paper, Bank of America.

Andersen, L., Sidenius, J., 2004. Extensions to the Gaussian copula: random recovery and random factor loadings. Journal of Credit Risk 1(1), 29–70.

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

Andersen, L., Sidenius, J., Basu, S., 2003. All your hedges in one basket. RISK 16, November 62–72.

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

Azizpour, S., Giesecke, K., 2008. Self-exciting corporate defaults: contagion vs. frailty. Working paper, Stanford University.

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

Collin-Dufresne, P., Goldstein, R. S., Helwege, J., 2003. Is credit event risk priced? Modeling contagion via the updating of beliefs. Working Paper, Haas School, University of California, Berkeley.

Cont R., Minca, A., 2007. Reconstructing portfolio default rates from CDO tranche spreads. Working paper, Columbia University.

Coval, J. D., Jurek, J. W., Stafford, E., 2008. Economic catastrophe bonds. Working paper, Havard University.

Cox J. C., Ingersoll, J. E., Ross, S. A., 1985. A theory of the term structure of interest rates. Econometrica 53, 129–151.

Crouhy, M., Jarrow, R. A., Turnbull, S. M., 2008. The subprime credit crisis of 07. Working paper, Natixis Funds, Cornell University, and University of Houston.

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

Demyanyk, Y., Hemert, O. V., 2008. Understanding the subprime mortgage crisis. Working paper, Federal Reserve Bank of St. Louis and New York University.

DeServigny, A., Renault, O., 2002. Default correlation: empirical evidence. Working paper, Standard and Poor's.

Duffie, D., 2007. Innovations in credit risk transfer: implications for financial stability. Working paper, Stanford University.

Duffie, D., Eckner, A., Horel, G., Saita, L., 2008. Frailty correlated default. Working paper, Stanford University, forthcoming Journal of Finance.

Duffie, D., Gârleanu, N., 2001. Risk and valuation of collateralized debt obligations. Financial Analysts Journal 57, 41–59.

Duffie, D., Pan, J., Singleton, K., 2000. Transform analysis and asset pricing for affine jump diffusions. Econometrica 68(6), 1343–1376.

Duffie, D., Saita, L., Wang, K., 2007. Multi-period corporate default prediction with stochastic covariates. Journal of Financial Economics 83, 635–665.

Duffie, D., Singleton, K., 1999. Modeling term structure of defaultable bonds. Review of Financial Studies 12, 687–720.

Duffie, D., Singleton, K.J., 2003. Credit Risk. Princeton University Press, Princeton, New Jersey.

Errais, E., Giesecke, K., Goldberg, L., 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., Kim, B., 2007. Estimating tranche spreads by loss process simulation. In Henderson, S. G., Biller, B., Hsieh, M.-H., Shortle, J., Tew, J. D., Barton, R. R.(Eds.), Proceedings of the 2007 Winter Simulation Conference, IEEE Press.

Glasserman, P., 2004. Monte Carlo Methods in Financial Engineering. Springer-Verlag, New York.

Hull, J., White, A., 2004. Valuation of a CDO and an n-th to default CDS without Monte Carlo simulation. Journal of Derivatives 12(2), 8–23.

Hull, J., White, A., 2007. Dynamic models of portfolio credit risk: a simplified approach. Working paper, University of Totonto.

Jarrow, R., Turnbull, S., 1995. Pricing options on financial securities subject to credit risk. Journal of Finance 50, 53–85.

Joe, H., 1997. Multivariate Models and Dependence Concepts. Chapman & Hall/CRC, New York.

Jorion, P., Zhang, G., 2007. Good and bad credit contagion: evidence from credit default swaps. Journal of Financial Economics 84(3), 860–883.

Joshi, M. S., Stacey, A. M., 2006. Intensity gamma: a new approach to pricing portfolio credit derivatives. RISK 19, July 78–83.

Lando, D., 1994. Three Essays on Contingent Claims Pricing. Ph.D. thesis, Cornell University.

Lando, D., 1998. On Cox processes and credit risky securities. Review of Derivatives Research 2, 99–120.

Lang, L., Stulz, R., 1992. Contagion and competitive intra-industry effects of bankruptcy announcements. Journal of Financial Economics 8, 45–60.

Lawrence, C. T., Zhou, J. L., Tits, A. L., 1997. User's Guide for CFSQP Version 2.5: a C Code for Solving (Large Scale) Constrained Nonlinear (Minimax) Optimization Problems. Software documentation, University of Maryland.

Li, D. X., 2000. On default correlation: a copula function approach. Journal of Fixed Income 9, 43–54.

Longstaff, F., Rajan, A., 2007. An empirical analysis of collateralized debt obligations. Working paper, University of California, Los Angeles.

Merton, R. C., 1974. On the pricing of corporate debt: The risk structure of interest rates. Journal of Finance 29, 449–470.

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

Papageorgiou, E., Sircar, R., 2007. Multiscale intensity models and name grouping for valuation of multi-name credit derivatives. Working paper, Princeton University.

Press, W. H., Teukolsky, S. A., Vetterling, W. T., Flannery, B. P., 2002. Numerical Recipes in C++: the Art of Scientific Computing, 2nd ed. Cambridge University Press, New York.

Schönbucher, P. J., 2003. Credit Derivatives Pricing Models: Models, Pricing and Implementation. John Wiley & Sons Ltd, Chichester, West Sussex, England.

Schönbucher, P. J., 1998. Term structure modelling of defaultable bonds. Review of Derivatives Research 2, 161–192.

Schönbucher, P. J., 2007. Time for a time-change: A new approach to multivariate intensity models of credit risk. Working paper, ETH Zürich.

United States Securities and Exchange Commission, July 2008. Summary report of issues identified in the commission staff's examinations of select credit rating agencies. Report, United States Securities and Exchange Commission.