An Overview of Credit Derivatives

Kay Giesecke*

Published in

Jahresbericht der Deutschen Mathematiker-Vereinigung,
Vol. **111**, No. 2, pp. 57-93, 2009

This version: October 24, 2011

Abstract

Credit risk is the distribution of financial loss due to a broken financial agreement, for example failure to pay interest or principal on a loan or bond. It pervades virtually all financial transactions, and therefore plays a significant role in financial markets. A credit derivative is a security that allows investors to transfer credit risk to other investors who are willing to take it. By facilitating the allocation of risk, these instruments have an important economic function. Yet they have hit the headlines recently. This paper gives an overview of credit derivatives. It discusses the mechanics of standard contracts, describes their application, and outlines the mathematical challenges associated with their analysis.

*Stanford University, Department of Management Science & Engineering, Stanford, CA 94305-4026, USA, Phone (650) 723 9265, Fax (650) 723 1614, email: giesecke@stanford.edu, web: www.stanford.edu/∼giesecke. I am grateful to Xiaowei Ding, Baeho Kim and Jack Kim for research assistance, Eymen Errais for providing market data, Alexander Schied for comments, and particularly Stefan Weber for stimulating discussions, comments and suggestions. I am also grateful for the suggestions of an anonymous referee.

1 Introduction

The financial crisis in the United States has repercussions on a global scale. Credit derivatives, especially collateralized debt obligations, have been pictured as one of the main culprits. This article provides an overview of credit derivatives. It discusses the mechanics of standard contracts, describes their application, and highlights the mathematical challenges associated with their analysis.

A credit derivative is a financial instrument whose cash flows are linked to the financial losses due to default in a pool of reference credit securities such as loans, mortgages, bonds issued by corporations or governments, or even other credit derivatives. The term "default" refers to an event that adversely affects the position of an investor in the reference securities. Examples include bankruptcy, failure to pay interest or principal according to schedule, debt moratorium, and restructuring of an issuer.

Credit derivatives facilitate the trading of credit risk, and therefore the allocation of risk among market participants. They resemble bilateral insurance contracts, with one party buying protection against default losses, and the other party selling that protection. This structure enables investors to take different sides and implement various investment and hedging strategies. For example, a fixed income investor may buy protection to hedge the default risk associated with a corporate bond position. An insurance company or hedge fund may act as the counterparty to this deal, and promise to pay potential default losses. The seller of protection speculates on the survival of the bond issuer, and gains investment exposure without having to commit the capital required to actually buy the bonds in the cash markets.

The growth in the volume of credit derivative transactions has exceeded expectations year by year after they were introduced in the early nineties. The trading of credit derivatives reached a peak in January 2008, when industry sources estimated the total notional of credit derivatives outstanding at 62 trillion dollars. During 2008, the near failure of investment bank Bear Stearns, the collapse of investment bank Lehman Brothers, insurance firm American International Group, and other market participants exposed the latent threats that credit derivatives can pose for the global financial system. The contract netting and unwinding motivated by these events and the development of a worldwide financial and economic crisis has pushed the outstanding contract notional to under 20 trillion dollars at the end of November 2008.

Governments and regulatory authorities call for more transparency and basic regulation of the credit derivatives market. While market liquidity has dried up, the economics behind a credit derivative contract remain sound. For example, banks making loans will continue to have a basic need for hedging loan exposures. Therefore, we expect the credit derivatives market to survive in a leaner, more transparent form that is subject to basic regulation. In this market, the risk management of exposures will likely have a more prominent role in the trading process than it currently has. The risk ratings of complex structured deals issued by rating agencies such as Moody's, Standard & Poor's or Fitch will likely be re-structured and improved. This development generates important research themes for economists, financial engineers, statisticians and mathematicians.

The primary objective of this article is to give a concise overview of the current credit derivative landscape. Section 2 illustrates some of the main features, application cases and potential problems of credit derivatives by providing an informal discussion of credit swaps, which reference a single credit obligation. Section 3 discusses a probabilistic model of default timing, on which our formal analysis in the subsequent sections is based. Section 4 analyzes corporate bonds, and then treats credit swaps in more detail. Section 5 discusses multi-name credit derivatives, which are referenced on a pool of issuers whose defaults are correlated.

2 Credit default swaps

To provide some perspective, this introductory section starts with an informal discussion of credit default swaps. Many more complex credit derivative instruments share the basic features of credit swaps.

A credit swap is a financial agreement that is individually negotiated between two investors. As indicated in Figure 1, it resembles an insurance contract. The protection seller compensates the protection buyer for the financial loss, if the issuer of a specified reference security defaults on his obligation before the maturity of the swap. The reference security can be a loan, a bond issued by a corporation or a sovereign nation, or a more complex instrument. The protection buyer obtains insurance against default of the reference security, and must pay a premium for this insurance. The premium is called the *credit swap spread*, is stated as a fraction of the swap notional per annum, and is typically paid quarterly until default or maturity, whichever is earlier.

Here is a basic example. A protection buyer purchases 5 year protection on an issuer with notional \$10 million at an annual swap spread of 300 basis points (one basis point corresponds to 0.01% of the notional). Suppose the reference issuer defaults 4 months after inception, and that the reference obligation has a recovery rate of 45%. This means that 55 cents on every dollar lent to the reference issuer are lost at default, an estimate typically obtained through a poll of market participants. Thus, 3 months after inception, the protection buyer makes the first spread payment, roughly equal to $\$10$ million $\times 0.03 \times 0.25 = \$75,000$.[1] At default, the protection seller compensates the buyer for the loss by paying $\$10$ million $\times(100\%-45\%) = \$5.5$ million, assuming the contract is settled in cash.[2] At the same time, the protection buyer pays to the seller the premium accrued since the last payment date, roughly equal to $\$10$ million $\times 0.03 \times 1/12 = \$25,000$. The payments

[1] In practice, the payment depends on the ratio of the actual number of days in the given quarter to the total number of days in the year, which is fixed at 360. We have approximated that ratio by 0.25.

[2] An alternative settlement convention is physical delivery. Here, the buyer delivers to the seller an asset of his choice from a specified pool of reference securities, in exchange for the notional. The terms of the swap may also stipulate a fixed cash payment at default.

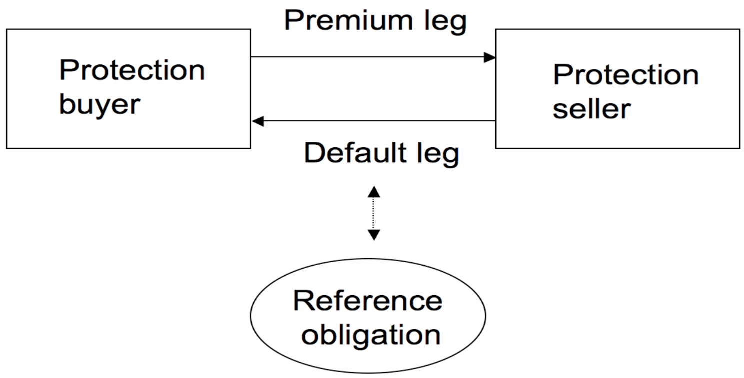

Figure 1: Mechanics of a credit default swap (CDS).

are netted. With these cash flows the swap expires; there are no further obligations in the contract.

The swap investors have opposing positions. The protection buyer has sold the credit risk associated with the reference obligation. If the buyer actually owns the reference security, then he has hedged the position against default. For a bank hedging its loans, this can lead to economic and regulatory capital relief, i.e. a decline in the capital required to support the bank's lending activities. If the buyer does not own the reference security, then he has entered into a speculative short position that seeks to benefit from a deterioration of the issuer's creditworthiness. If the target scenario unfolds, the buyer simply sells the contract to another party, for a profit because default protection commands a higher premium after a deterioration of the issuer's credit quality. The possibility of buying default insurance on a bond without actually owning the bond distinguishes the credit swap from a classical insurance contract.

The protection seller has bought the credit risk of the reference security. Relative to buying the reference security directly, the swap position has the advantage of not requiring a capital outlay at inception. To make this clear, think of an investor who wishes to collect compensation for assuming the credit risk of a given corporation. The investor could purchase a coupon bond issued by the corporation by paying the face value (or current price) of the bond, and then collecting the coupons paid by the issuer. The investor could alternatively assume exposure to the issuer's credit risk by selling protection in a credit swap referenced on the issuer's bond. This would involve only a commitment to compensate the swap protection buyer for the potential losses due to default of the issuer before the swap maturity, but no initial cash flow. In return for this commitment, the seller collects the swap spread from the swap protection buyer.

Since the credit swap is unfunded, it does not appear as a liability on the balance sheet of the swap protection seller. This "off-balance sheet" nature of the contract is a main feature of many credit derivatives. While attractive to many investors, this feature can also be liability for the financial system. This is because an investor can act as a protection seller in a large number of contracts without disclosing these deals, and accumulate a tremendous exposure to defaults that is practically “invisible” to regulators and other market participants. Although a protection seller’s economic position is effectively that of an insurance firm, the seller is not subject to the insurance industry’s regulation. Insurance regulation imposes capital requirements compatible with the insurance risks. An insurer’s capital limits the policies a firm can write, and therefore the total risk that can be taken. Without transparency and regulatory capital constraints, the total exposure of a protection seller is not limited, and few defaults can exhaust the capital of an over-exposed investor. When an investor is wiped out, the remaining swap contracts become null, leaving the corresponding protection buyers exposed to further defaults. A chain reaction might ensue, jeopardizing the stability of the financial system. Recently, such a meltdown became a real possibility with the near-default of Bear Stearns, an investment bank that had written protection in a large number of deals. A disaster could eventually be avoided since JP Morgan, another bank, bought Bear Stearns and entered into its contracts, backed by guarantees from the US government. The government chose not to provide similar support to the investment bank Lehman Brothers. Lehman had sold protection on a large number of firms and was itself a reference entity to countless other contracts, and its default brought the system to a near breakdown. The collapse of American International Group, an insurance firm that had written protection on a large scale as part of its investment strategy, was only averted by massive government intervention. These events call for meaningful regulation of the credit swap and general credit derivatives market, an effort that governments, central banks and regulatory authorities actively discuss.

The basic swap pricing problem is to determine the fair swap spread at inception. The swap spread must equate the present value at inception of the premium payments and the present value of the payments at default. After inception, the swap must be marked to the market, i.e. its current market value must be determined. For the protection seller, the mark-to-market value is given by the difference of the present values of premium and default payments evaluated at the prevailing market conditions.[3] If, for example, the credit risk of the reference issuer has deteriorated after inception, then the fixed swap spread does not provide sufficient compensation for the default risk anymore. The protection seller could obtain a higher spread at the current conditions. Therefore, the seller’s mark-to-market value may be negative in this scenario. Of course, the loss of the seller is the gain of the protection buyer, whose mark-to-market value is the negative of the seller’s mark-to-market value. If the buyer has implemented a short credit position as described above, then in this scenario the buyer could sell his position, with a profit equal to the current mark-to-maket value, less the value of the premia already paid.

In a liquid swap market, the swap rates provide a sensitive measure of an issuer’s creditworthiness, from which we can extract issuer default probabilities for multiple horizons. These market-implied probabilities reflect investors’ forward-looking expectations about a

[3]Marking a position to the market may be difficult when the market is illiquid and there are few buyers or sellers of default risk.

firm's credit quality. They can be used to value more complex, less liquid credit derivative instruments, and to design credit investment strategies. To facilitate these applications, we require a probabilistic model of default timing, to which we turn next.

3 Default point process

The uncertainty in the economy is modeled by a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is a set that represents the possible states of the world, $\mathcal{F}$ is a sigma-field on Ω, and $\mathbb{P}$ is a probability measure on $\mathcal{F}$. The flow of information accessible to investors is modeled by a filtration $\mathbb{F}$, i.e. an increasing family $(\mathcal{F}_t)_{t\geq 0}$ of sub-sigma-fields of $\mathcal{F}$. Intuitively, $\mathcal{F}_t$ represents the events observable at t. We assume that $\mathbb{F}$ is right-continuous and contains all $\mathbb{P}$-null sets in $\mathcal{F}$; see Dellacherie & Meyer (1982).

A firm's default time is modeled by a stopping time S, i.e., a non-negative random variable with the property that $\{S > t\} \in \mathcal{F}_t$. We consider a sequence of stopping times T^k with $T^0 = 0$ that strictly increases to ∞, almost surely. The T^k represent the ordered default times in a pool of firms. They generate a counting process N given by

$$N_t = \sum_{k\geq 1} 1_{\{T^k \leq t\}}. \tag{1}$$

Here and below, 1_B is the indicator function of an event $B \in \mathcal{F}$. That is, $1_B(\omega)$ is equal to one if $\omega \in B$ and zero otherwise.

A credit derivative is a security with cash flows that depend on the value of the process N at a set of times. For example, the contract may stipulate the payment of N_T at a maturity date T. In this case, the buyer of the derivative is protected against the defaults in the pool, assuming the loss at a default is less than unity.

To determine the value of a credit derivative, we require the distribution of N. Below, we specialize into the approach of Giesecke & Zhu (2010) to provide a formula for the Laplace transform of N. The distribution of N, obtained by transform inversion, can be used to calculate the value of a derivative with payoff $g(N_T)$ at maturity T, where g is an integrable function on $\{0, 1, 2, \ldots\}$ specified by the derivative terms. We illustrate this in Sections 4 and 5. Section 3.3 comments on alternative computational methods for the valuation of credit derivatives whose cash flows take a complicated form.

The development of the transform formula requires some prerequisits from the theory of stochastic processes, which is treated in Dellacherie & Meyer (1982), Karatzas & Shreve (1988), Protter (2004) and others. Readers that are not interested in the technical details but rather in the structure of standard credit derivatives can skip this section and proceed directly to Sections 4 and 5.

3.1 Transform analysis

Since the default process N has increasing sample paths, the Doob-Meyer decomposition theorem guarantees the existence of an increasing, predictable process A starting at 0 such that $M = N - A$ is a local martingale. The process A counteracts the increasing tendency of the default process; it is called the compensator to N. The compensator depends on the underlying filtration $\mathbb{F}$ and the reference measure $\mathbb{P}$; it is transformed with a change of $\mathbb{F}$ or $\mathbb{P}$. The analytic properties of A correspond to probabilistic properties of the default stopping times. The sample paths of A are continuous if and only if the T^k are *totally inaccessible.* Totally inaccessible stopping times formalize the intuitive concept of random times that come unannounced and take investors by surprise. There is empirical evidence in support of this property; see Sarig & Warga (1989), for example.

Assume the default times are totally inaccessible and that $\exp(A_T)$ is integrable for a fixed horizon $T > 0$. We use the measure change argument of Giesecke & Zhu (2010) to develop a formula for the conditional Laplace transform

$$\varphi(u, t, T) = \mathbb{E}(e^{-u(N_T - N_t)} \,|\, \mathcal{F}_t), \quad u \geq 0. \tag{2}$$

The corresponding conditional distribution can be recovered from the relation

$$\mathbb{P}(N_T - N_t = k \,|\, \mathcal{F}_t) = \lim_{v \downarrow 0} \frac{1}{k!} \partial_v^k \varphi(-\log v, t, T), \quad k = 0, 1, 2, \ldots \tag{3}$$

Alternatively, we can extend (2) to the complex plane and invert the Fourier transform so obtained: $\mathbb{P}(N_T - N_t = k \,|\, \mathcal{F}_t) = (2\pi)^{-1} \int_{-\pi}^{\pi} \varphi(-iv, t, T) \exp(-ikv) dv$. This formulation facilitates the application of the Fast Fourier Transform, which may have computational advantages over formula (3), depending on the particular setting.

For $u \geq 0$ define the process $Z(u)$ by

$$Z_t(u) = \exp(\psi(u) A_t - u N_t), \qquad \psi(u) = 1 - \exp(-u). \tag{4}$$

Since A is continuous, an application of Stieltjes integration by parts on the product $Z_t(u) = \exp(\psi(u) A_t) \exp(-u N_t)$ leads to the alternative expression

$$Z_t(u) = 1 - \psi(u) \int_0^t Z_{s-}(u) dM_s,$$

showing that $Z(u)$ is the stochastic exponential of the scaled local martingale $-\psi(u)M$. Hence, $Z(u)$ is a local martingale itself. The integrability condition on the compensator guarantees that the stopped process $Z^T(u)$ is a uniformly integrable martingale.

The family of martingales $Z(u)$ indexed by u induces a family of equivalent probability measures $\mathbb{P}^u$ on $\mathcal{F}_T$ by $\mathbb{P}^u(B) = \mathbb{E}(Z_T(u) 1_B)$ for $B \in \mathcal{F}_T$. Each measure $\mathbb{P}^u$ corresponds to a conditional Laplace transform of the compensator,

$$\mathcal{L}^u(v, t, T) = \mathbb{E}^u(e^{-v(A_T - A_t)} \,|\, \mathcal{F}_t), \quad u, v \geq 0. \tag{5}$$

The $\mathbb{P}$-Laplace transform (2) can be expressed in terms of the $\mathbb{P}^u$-Laplace transform (5),

$$\begin{aligned}\varphi(u,t,T) &= \mathbb{E}(e^{-\psi(u)(A_T-A_t)} Z_T(u)/Z_t(u) \,|\, \mathcal{F}_t) \\ &= \mathbb{E}^u(e^{-\psi(u)(A_T-A_t)} \,|\, \mathcal{F}_t) \\ &= \mathcal{L}^u(\psi(u),t,T). \end{aligned} \tag{6}$$

Formula (6) states that the Laplace transform of a counting process with totally inaccessible arrivals is given by the Laplace transform of the compensator evaluated at the Laplace exponent ψ of the standard Poisson process. The Laplace transform of A is taken under an equivalent measure $\mathbb{P}^u$. To calculate that transform, consider the specification of A in terms of a non-negative *intensity* process λ such that

$$A_t = \int_0^t \lambda_s \, ds \tag{7}$$

almost surely. The process λ is the density of the random measure associated with A relative to Lebesgue measure. It can be interpreted as the conditional event arrival rate, in the sense that $\mathbb{E}(N_{t+\Delta} - N_t \,|\, \mathcal{F}_t) \approx \lambda_t \Delta$ for "small" $\Delta > 0$, see Brémaud (1980). With the representation (7), the Laplace transform (5) takes the form

$$\mathcal{L}^u(v,t,T) = \mathbb{E}^u(e^{-v\int_t^T \lambda_s ds} \,|\, \mathcal{F}_t). \tag{8}$$

The expectation on the right hand side of equation (8) is a familiar expression in finance; see Duffie (1996), for example. Consider a government bond that pays 1 at time T. The price of this bond at $t \leq T$ is analogous to formula (8) once we regard the process $v\lambda$ as the short-term rate of interest. The explicit calculation of the bond price is well understood for a large class of interest rate processes, see Duffie, Pan & Singleton (2000) and Leippold & Wu (2002). This observation motivates us to adopt a parametric model formulation from this bond pricing literature for the purpose of specifying the intensity λ. However, before we can apply the results from the bond pricing literature, we need to understand how the dynamics of the intensity λ are adjusted when the measure is changed to $\mathbb{P}^u$. This is necessary since the dynamics of λ are typically specified under $\mathbb{P}$, while the Laplace transform (8) is taken under $\mathbb{P}^u$. To discuss the adjustment of the dynamics of λ, let V be a local martingale relative to the reference measure $\mathbb{P}$. Girsanov's theorem as stated in Dellacherie & Meyer (1982) implies that the process

$$V + \psi(u)\langle V, N\rangle \tag{9}$$

is a $\mathbb{P}^u$-local martingale. Here, $\langle V, N\rangle$ is the $\mathbb{P}$-conditional covariation, i.e. the compensator of the quadratic variation $[V, N]$ relative to $\mathbb{P}$. Consider a local martingale V that does not have jumps in common with N. Here $[V, N] = 0$, and therefore $\langle V, N\rangle = 0$. Thus, V remains a local martingale under $\mathbb{P}^u$. Now consider the local martingale $V = N - A$, whose jumps coincide with those of N. The quadratic variation $[V,N] = [N,N] = N$, whose compensator is $\langle V,N\rangle = A$. Thus, from (9),

$$V + \psi(u)\langle V,N\rangle = V + \psi(u)A = N - e^{-u}A$$

is a local martingale under $\mathbb{P}^u$, meaning N has $\mathbb{P}^u$-intensity $e^{-u}\lambda$. The measure change calls for a deterministic scaling of the intensity that depends on the variable u. In Section 3.2 below, we show how these observations can be applied to calculate the $\mathbb{P}^u$-dynamics of λ for a classical specification of λ under $\mathbb{P}$ adopted from the bond pricing literature.

The measure change underlying the transform formula (6) is not necessary when N is a doubly-stochastic Poisson process. The process N is called a doubly-stochastic Poisson process if the compensator A is adapted to a right-continuous and complete sub-filtration $\mathbb{G} = (\mathcal{G}_t)_{t\geq 0}$ of $\mathbb{F}$, and the conditional distribution of $N_T - N_t$ given $\mathcal{F}_t \vee \mathcal{G}_T$ is the Poisson distribution with parameter $A_T - A_t$, see Brémaud (1980). More precisely, for all $t < T$ and $u \geq 0$ we have

$$\mathbb{E}(e^{-u(N_T-N_t)} \,|\, \mathcal{F}_t \vee \mathcal{G}_T) = e^{-\psi(u)(A_T-A_t)}. \tag{10}$$

Taking conditional expectations relative to $(\mathbb{P}, \mathcal{F}_t)$ on both sides of equation (10), we get a formula for the Laplace transform of a doubly-stochastic Poisson process:

$$\begin{aligned}\varphi(u,t,T) &= \mathbb{E}(e^{-\psi(u)(A_T-A_t)} \,|\, \mathcal{F}_t) \\ &= \mathcal{L}^0(\psi(u),t,T)\end{aligned} \tag{11}$$

where $\mathcal{L}^0(v,t,T)$ is the Laplace transform of the compensator under $\mathbb{P} = \mathbb{P}^0$. Compare formula (11) with the general transform formula (6). In the general, non doubly-stochastic case, the measure change leads to a transform formula that has the same structure as the doubly-stochastic formula (11). In this sense, the measure change preserves the doubly-stochastic setting. The special feature of this doubly-stochastic setting is that the arrivals of N are not allowed to influence the dynamics of A. This means that a doubly-stochastic Poisson process cannot be self-exciting: an arrival of the process cannot increase the likelihood of further arrivals. This is because the doubly-stochastic compensator is adapted to $\mathbb{G}$, and the event $\{N_t = k\}$ can never be contained in $\mathcal{G}_t$. Since the change of measure to $\mathbb{P}^u$ "absorbs" any feedback from N to the dynamics of A, we call $\mathbb{P}^u$ the *correlation-neutral probability measure.*

The measure change argument underlying the transform formula (6) is analogous to a similar argument proposed by Carr & Wu (2004) to calculate the Fourier transform of a time-changed Lévy process. Carr & Wu (2004) suggest that this transform is given by the Laplace transform of the time change evaluated at the characteristic exponent of the Lévy process. The Laplace transform is calculated under a complex measure defined by the time-changed Wald martingale associated with the Lévy process. This argument leads to a formula for the Fourier transform of a counting process N if N is realized as a time-changed Poisson process. In this case the compensator A takes the role of the time change and Z coincides with the time-changed Wald martingale.

The transform formula (6) can be extended in several directions. As detailed in Giesecke & Zhu (2010), an analogous formula can be developed for a vector of correlated point processes with random jump sizes, incorporating the effects of discounting at a stochastic interest rate and of a random cash flow at a horizon.

3.2 Affine point processes

The transform formula (6) is of practical use whenever the $\mathbb{P}^u$-Laplace transform (5) of the point process compensator A can be explicitly calculated. In this section we discuss an example. We consider a default point process N whose intensity is driven by an affine jump diffusion in the sense of Duffie et al. (2000). The resulting transform formula is a special case of the transform formula for general affine jump diffusions developed in Duffie et al. (2000) and Duffie, Filipovic & Schachermayer (2003). Our measure change argument provides an alternative route to the formula for a finite-activity affine jump process.

Suppose the observation filtration $\mathbb{F}$ is generated by a Markov process X that is a strong solution to the stochastic differential equation

$$dX_t = \mu(X_t)\, dt + \sigma(X_t)\, dW_t + \delta\, dJ_t, \qquad X_0 \in \mathbb{R}. \tag{12}$$

Here W is a standard Brownian motion, $\mu(X)$ is the drift process, $\sigma(X)$ is the volatility process, $\delta \geq 0$ is a sensitivity parameter and $J_t = \sum_{k=0}^{N_t} J^k$ is a point process whose jump times are those of the default counting process N. The random jump sizes J^k are drawn, independently of one another, from a distribution ν on $\mathbb{R}_+$ that has no mass at zero.

We interpret X as a risk factor process that describes the state of the economy. Defaults influence the economic state in that X jumps at each event. The size of the impact is random, and the sensitivity is controlled by δ. Suppose the compensator A of N has intensity $\lambda = \Lambda(X)$ for some non-negative function Λ on $\mathbb{R}$. Then, there is feedback from N to its intensity via X. The default point process N is self-exciting, in the sense that an event tends to increase the likelihood of further events.

By formula (6), the calculation of the Laplace transform of N reduces to the calculation of the Laplace transform (5) of A under the correlation-neutral measure. To obtain an explicit expression for the transform (5), we impose an additional structure on X and λ. We assume that X and the functions μ, σ^2 and Λ satisfy

$$\mu(x) = K_0 + K_1 x, \quad \sigma(x)^2 = H_0 + H_1 x, \quad \Lambda(x) = \Lambda_0 + \Lambda_1 x,$$

for constant coefficients such that $\Lambda(X)$ is non-negative and $\exp(\int_0^T \Lambda(X_s) ds)$ is integrable for the horizon $T > 0$. Under these assumptions, the state process X is an affine jump diffusion in the sense of Duffie et al. (2000), and the counting process N is called an affine point process. Special cases of N include the Poisson process ($H_0 = H_1 = \delta = 0$), the birth process ($K_0 = K_1 = H_0 = H_1 = 0$), and the Hawkes process ($H_0 = H_1 = 0$). Proposition 1 in Duffie et al. (2000) states technical conditions such that

$$\mathbb{E}(e^{-\int_t^T R(X_s)ds+zX_T} \mid \mathcal{F}_t) = e^{a(t)+b(t)X_t} \tag{13}$$

where $R(x) = \rho_0+\rho_1 x$ for constants ρ_0 and ρ_1, and the coefficient functions $b(t) = b(z,t,T)$ and $a(t) = a(z,t,T)$ satisfy the ordinary differential equations

$$\partial_t b(t) = \rho_1 - K_1 b(t) - \frac{1}{2}H_1 b(t)^2 - \Lambda_1(\theta(\delta b(t)) - 1) \tag{14}$$

$$\partial_t a(t) = \rho_0 - K_0 b(t) - \frac{1}{2}H_0 b(t)^2 - \Lambda_0(\theta(\delta b(t)) - 1) \tag{15}$$

with boundary conditions $b(T) = z$ and $a(T) = 0$ and jump transform

$$\theta(c) = \int_{\mathbb{R}_+} e^{cz} d\nu(z), \quad c \in \mathbb{R}, \tag{16}$$

whenever this integral is well-defined. For $\rho_0 = \rho_1 = 0$, formula (13) yields the transform of the compound birth process $X = \delta J$ (set $K_0 = K_1 = H_0 = H_1 = 0$ in (14)–(15)). To get the transform for more general point processes whose intensity is driven by X, we need to construct an enlarged state vector (X, N) by appending the counting process N to the original state X, see Errais, Giesecke & Goldberg (2010). Proposition 1 in Duffie et al. (2000) is then applied to (X, N). Rather than following this route, we show how to use our transform formula (6) in the affine setting.

Formula (6) requires the Laplace transform (5) relative to the correlation-neutral measure $\mathbb{P}^u$. We can apply Proposition 1 of Duffie et al. (2000) to obtain a formula for this Laplace transform. To this end, note that the $\mathbb{P}^u$-dynamics of the state X are described by equation (12), where W is a $\mathbb{P}^u$-Brownian motion, $J_t = \sum_{k=0}^{N_t} J^k$ is a point process whose jump times arrive with $\mathbb{P}^u$-intensity

$$e^{-u}\Lambda(X) = e^{-u}\Lambda_0 + e^{-u}\Lambda_1 X,$$

and where the J^k are drawn, independently of one another, from the distribution ν relative to $\mathbb{P}^u$. Here, we use the fact that W has no jumps in common with N and therefore remains a Brownian motion under $\mathbb{P}^u$, that the distribution of the J^k is not affected by the measure change, and that the compensator of N is scaled by e^{-u} when the measure is changed. Now we apply Proposition 1 in Duffie et al. (2000), taking $\mathbb{P}^u$ as reference measure, to get

$$\mathcal{L}^u(v,t,T) = \mathbb{E}^u(e^{-v\int_t^T \Lambda(X_s)ds} \mid \mathcal{F}_t) = e^{\alpha(t)+\beta(t)X_t} \tag{17}$$

where the coefficient functions $\beta(t) = \beta(u,v,t,T)$ and $\alpha(t) = \alpha(u,v,t,T)$ satisfy the ordinary differential equations

$$\partial_t \beta(t) = v\Lambda_1 - K_1\beta(t) - \frac{1}{2}H_1\beta(t)^2 - e^{-u}\Lambda_1(\theta(\delta\beta(t)) - 1) \tag{18}$$

$$\partial_t \alpha(t) = v\Lambda_0 - K_0\beta(t) - \frac{1}{2}H_0\beta(t)^2 - e^{-u}\Lambda_0(\theta(\delta\beta(t)) - 1) \tag{19}$$

with boundary conditions $\beta(T) = \alpha(T) = 0$ and jump transform (16). Thus, the conditional Laplace transform (2) of an affine point process N is given in terms of the solutions to the equations (18)–(19) as

$$\varphi(u,t,T) = \mathcal{L}^u(\psi(u),t,T) = \exp(\alpha(u,\psi(u),t,T) + \beta(u,\psi(u),t,T)X_t). \tag{20}$$

Compare equations (17)–(19) with equations (13)–(15). The measure change influences only the jump terms in the ODEs (18)–(19). It has no effect when $\delta = 0$ in the risk factor dynamics (12). In this case, N is doubly-stochastic. There is no feedback from N to its intensity, and the $\mathbb{P}^u$-Laplace transform (8) agrees with its counterpart under $\mathbb{P}$.

3.3 Credit derivatives valuation

There are two alternative formulations of the credit derivatives valuation problem:

(1) Model the default point process N under an equivalent martingale measure, relative to which the discounted price process of any traded security is a martingale. Apply the transform formula (6) with the martingale measure compensator of N.

(2) Model N under the actual probability measure that represents the empirical likelihood of default events. Specify a martingale measure equivalent to the physical measure, and calculate the martingale measure compensator from the physical measure compensator by Girsanov's theorem. Apply (6).

In Sections 4 and 5 below, we follow the first approach, and specify the "risk-neutral" behavior of N under $\mathbb{P}$, which we take to be an equivalent martingale measure relative to a constant risk-free interest rate r. To exemplify this approach, fix a parametric model for the $\mathbb{P}$-compensator A, for instance the affine specification analyzed in Section 3.2 above. Consider a credit derivative with payoff $g(N_T)$ at maturity T, where g is an integrable function on $\{0, 1, 2, \ldots\}$ specified by the terms of the derivative contract. By the martingale property of discounted prices under $\mathbb{P}$, an arbitrage-free value of the derivative at time $t \leq T$ is given by the conditional expectation

$$e^{-r(T-t)}\mathbb{E}(g(N_T) \,|\, \mathcal{F}_t) = e^{-r(T-t)} \sum_{k \geq 0} g(N_t + k)\mathbb{P}(N_T - N_t = k \,|\, \mathcal{F}_t) \tag{21}$$

where the conditional probabilities $\mathbb{P}(N_T - N_t = k \,|\, \mathcal{F}_t)$ are obtained from the Laplace transform (2) of N via the inversion formula (3). If default arrivals are totally inaccessible and $\exp(A_T)$ is integrable, then the Laplace transform (2) is given by the Laplace transform of A with respect to the correlation-neutral measure, see formula (6).

In practice, the parameters of the model A are chosen through a calibration procedure. Here (21) is applied to a set of marketed credit derivatives, i.e., different payoff functions g available for trading. The parameters are chosen so as to minimize the discrepancy between the prices of the marketed derivatives and the prices generated by the model for these securities. In practice, the calibrated model is often used to assign prices to non-traded "exotic" credit derivatives. This approach, however, obscures the fact that the credit derivatives market is typically incomplete. That is, not every credit derivative can be perfectly hedged by dynamic trading in the marketed securities. A derivative carries intrinsic risk that cannot be hedged away. In this situation, the practice of pricing exotic derivatives via (21) does not systematically account for the costs of hedging or the residual risks. The pricing rule (21) assigns to the exotic derivative a price that pretends that the security can be perfectly hedged. The proper quantification and valuation of the residual risk motivates the design of incomplete market models. We refer to Föllmer & Schied (2004) for a comprehensive treatment of this topic and Staum (2008) for a survey paper.

The second approach outlined above is more complete than the first approach, and would be used for empirical time-series applications that require the distribution of N under the actual measure and the specification and estimation of the risk premia parameters that determine the change of measure from actual to risk-neutral probabilities, as in Berndt, Douglas, Duffie, Ferguson & Schranz (2005), Eckner (2007), and Azizpour, Giesecke & Kim (2011).

3.4 Monte Carlo simulation

An alternative to the transform approach to credit derivatives valuation illustrated above is Monte Carlo (MC) simulation. MC methods have a broader scope than transform methods. They can be used, for example, to treat a credit derivative with path-dependent payoff $f(t,(N_s)_{0\leq s\leq t},\ldots)$. The basic idea is to estimate the expected payoff by its sample average obtained from a large number of independent trials.

To generate a path of the counting process N, we sequentially generate the event times T_n, starting from $T_0=0$. There are several methods for the generation of the T_n. If the inter-arrival intensity is bounded almost surely, then one can draw a sample of T_n using the thinning scheme of Lewis & Shedler (1979). Ogata (1981) and Glasserman & Merener (2003) establish the validity of this scheme when the inter-arrival intensity is a bounded stochastic process rather than a deterministic function of time as assumed by Lewis & Shedler (1979). In this scheme, one generates an exponential random variable $\mathcal{E}_1$ with parameter H, and accepts $T_{n-1}+\mathcal{E}_1$ as T_n with probability $\nu_{\mathcal{E}_1}/H$, where ν_t is the value of the inter-arrival intensity at t. In the case of rejection, one iterates by redefining the bound and by drawing another exponential $\mathcal{E}_2$, now starting from $T_{n-1}+\mathcal{E}_1$. This scheme is applicable whenever the inter-arrival intensity is bounded almost surely, and one can draw samples from its distribution at candidate arrival times. If these samples are exact, then so are the samples of the T_n. The methods of Beskos & Roberts (2005) and Chen (2009) can be used to generate exact samples of a broad range of diffusion inter-arrival intensity processes. The method of Giesecke & Smelov (2010) supplies exact samples of virtually any jump-diffusion process. These exact methods eliminate the need to discretize the inter-arrival intensity process.

The boundedness of the inter-arrival intensity is a serious restriction. It fails for many standard model specifications, limiting the applicability of the thinning scheme. Examples include the case that the intensity follows a Feller diffusion, a more general affine jump-diffusion as in Section 3.2 above, or an exponential Ornstein-Uhlenbeck process. In these and many other cases, the probability of the inter-arrival intensity hitting any finite level H in finite time is strictly positive. Giesecke, Kim & Zhu (2010) extend the thinning scheme to the case where the inter-arrival intensity is unbounded.

There are alternative methods to simulate N when the inter-arrival intensity is unbounded. The most widely applicable scheme is a discretization scheme based on the time change theorem of Meyer (1971). If the compensator A is continuous and increases to ∞ almost surely, then the default process N is a standard Poisson process under a change of time defined by A, relative to the time-changed filtration. This implies that, for a sequence $(\mathcal{E}_n)$ of i.i.d. standard exponential variables,

$$T_n \stackrel{d}{=} \inf\left\{t \geq 0 : A_t \geq \mathcal{E}_1 + \cdots + \mathcal{E}_n\right\}. \tag{22}$$

We exploit this relation by approximating T_n by $\widehat{T}_n = \inf\{t \geq 0 : \widehat{A}_t \geq \mathcal{E}_1 + \cdots + \mathcal{E}_n\}$ for $n \geq 1$, taking $\widehat{T}_0 = T_0 = 0$. Here, $\widehat{A}$ approximates the compensator A on a discrete-time grid. A particular choice is $\widehat{A}_t = \int_0^t V_s ds$, where

$$V_t = \widehat{\lambda}_{k(t)} + (\widehat{\lambda}_{\lceil t/h \rceil h} - \widehat{\lambda}_{k(t)}) \frac{t - k(t)}{\lceil t/h \rceil h - k(t)}, \tag{23}$$

where $k(t) = \max(\lfloor t/h \rfloor h, \widehat{T}_{\widehat{N}_t})$, h is the time step, $\widehat{N}$ is the counting process of the $\widehat{T}_n$, and $\widehat{\lambda}$ is the approximation of λ. Other choices include $V_t = \widehat{\lambda}_{k(t)}$ and $V_t = \widehat{\lambda}_{\lceil t/h \rceil h}$. One may also consider a "fixed-grid scheme," in which default times are allowed to occur only at the discrete-time grid points. Here, A is approximated by a piece-wise constant function given by the left or right Riemann sums or the trapezoidal rule. Barring special cases (such as the affine model of Section 3.2), the values of the approximation $\widehat{\lambda}$ at the grid and other points must be generated by a discretization method.

The time-scaling scheme has many advantages, including the relative ease of implementation and the wide applicability. On the other hand, it generates a biased simulation estimator. The bias is due to the approximation of the compensator, whose paths can almost never be simulated exactly. The magnitude of the bias is hard to quantify. This makes it difficult to obtain valid confidence intervals for the simulation estimator. It also makes it difficult to determine the optimal allocation of the computational budget between the number of trials and the number of time steps.

Alternative exact simulation methods seek to avoid the discretization of the compensator. The thinning method described above is an example. Another is the inverse transform scheme. This scheme is applicable whenever the $\mathcal{F}_{T_{n-1}}$-conditional distribution of an inter-arrival time $T_n - T_{n-1}$ is computationally tractable, and one can generate a sample of the state variables at T_{n-1}. In the affine setting of Section 3.2, the inter-arrival intensity follows a Feller diffusion process so $\mathbb{P}(T_n - T_{n-1} > t \,|\, \mathcal{F}_{T_{n-1}})$ is an exponentially affine function of $\lambda_{T_{n-1}}$ with coefficients provided by Cox, Ingersoll & Ross (1985). The value of $\lambda_{T_{n-1}}$ is sampled from the appropriate conditional density using an acceptance/rejection scheme. Giesecke & Kim (2007) provide the details.

Another exact scheme is proposed by Giesecke, Kakavand & Mousavi (2010). Their method is based on a filtering argument. The point process is first projected onto its own filtration and then sampled in this coarser subfiltration. The sampling is based on the subfiltration-intensity, which is deterministic between event times and therefore facilitates the use of exact schemes, including the thinning algorithm.

Giesecke, Kakavand, Mousavi & Takada (2010) provide a method for the exact sampling of N_T at a fixed time T. The idea is to construct a continuous time Markov chain with state-space $\{0, 1, 2, \ldots\}$ whose value at T has the same distribution as N_T. It can be shown that such a mimicking chain exists, and that its transition rate is given by a conditional expectation of λ that can be computed for many standard models of λ. The construction reduces the original MC problem to one involving a simple Markov chain, which can be sampled exactly (using, once again, the thinning algorithm). This scheme leads to an unbiased simulation estimator of the price of a derivative whose payoff is an arbitrary function of N_T. The exact schemes outlined above lead to unbiased estimators of the prices of derivatives with payoffs that may be path-dependent.

4 Single-name credit derivatives

This section discusses the valuation of credit derivatives referenced on a given firm. The issuer's default time is taken to be the first jump time T^1 of the counting process N. The corresponding default process is $N^1 = \min(N, 1)$, the process N stopped at its first jump. The financial loss at default is modeled by an $\mathcal{F}_{T^1}$-measurable random variable ℓ^1, which is independent of T^1 and has expectation $\ell = \mathbb{E}(\ell^1)$. At any time $t < T^1 \wedge T = \min(T, T^1)$, the firm's risk-neutral conditional survival probability satisfies almost surely

$$\mathbb{P}(T^1 > T \,|\, \mathcal{F}_t) = \mathbb{P}(N_T = 0 \,|\, \mathcal{F}_t) = \lim_{u \uparrow \infty} \varphi(u, t, T) = \lim_{u \uparrow \infty} \mathcal{L}^u(\psi(u), t, T) \tag{24}$$

where the last equality is valid if default arrivals are totally inaccessible and $\exp(A_T)$ is integrable, see Section 3.1 above. If, for example, N is an affine point process as in Section 3.2, then $\mathcal{L}^u(\psi(u), t, T)$ is given by formula (20), and with the corresponding coefficient equations (18) and (19), for $t < T^1 \wedge T$ we then get the formula

$$\mathbb{P}(T^1 > T \,|\, \mathcal{F}_t) = \exp(a(t, T) + b(t, T)X_t) \tag{25}$$

where the coefficient functions $b(t) = b(t, T)$ and $a(t) = a(t, T)$ satisfy the ODEs

$$\partial_t b(t) = \Lambda_1 - K_1 b(t) - \frac{1}{2} H_1 b(t)^2 \tag{26}$$

$$\partial_t a(t) = \Lambda_0 - K_0 b(t) - \frac{1}{2} H_0 b(t)^2 \tag{27}$$

with boundary conditions $b(T) = a(T) = 0$. Equations (26) and (27) indicate that the feedback from N to its intensity, which comes from the jump term J in the state equation (12), plays no role for the survival probability (25). Intuitively, the conditional distribution of T^1 is governed by the dynamics of the intensity before T^1. The behavior of the intensity at the event time is not relevant for this distribution.

4.1 Corporate zero-coupon bond

Corporate bonds are issued by firms to raise funds. While they are not credit derivatives in a strict sense, we use them as basic building blocks for credit derivatives. A zero coupon bond with unit face value and maturity T pays

- The face value 1 at $T < T^1$,
- The recovery $(1 - \ell^1)$ of face value at $T^1 \leq T$.

Since $\mathbb{P}$ is an equivalent martingale measure, the value of the *survival cash flow* at a time t prior to default is given by the discounted $\mathbb{P}$-expected payoff

$$\mathbb{E}(F(t,T)(1 - N_T^1) \,|\, \mathcal{F}_t) = F(t,T)\mathbb{P}(T^1 > T \,|\, \mathcal{F}_t), \tag{28}$$

where $F(t,T) = \exp(-r(T-t))$ is the price at time t of a unit face value, T-maturity zero coupon government bond without default risk. The corporate bond is worth less than the government bond to account for the issuer's risk of default.

If N is doubly-stochastic with intensity λ, then the value of the survival cash flow (28) takes the form $\mathbb{E}(e^{-\int_t^T (r+\lambda_s)ds} \,|\, \mathcal{F}_t)$, which can be interpreted as the value of the certain cash flow 1 when the discount rate is $r + \lambda$ rather than the risk-free rate r. Thus, the intensity can be interpreted as a short-term credit spread that compensates the bond investor for assuming the issuer's risk of default over an infinitesimal period.

The value of the *recovery cash flow* at a time t prior to default is

$$\mathbb{E}(F(t,T^1)(1 - \ell^1)N_T^1 \,|\, \mathcal{F}_t) = (1 - \ell)R_t(T) \tag{29}$$

where $R_t(T)$ is the pre-default value of a unit recovery payment at $T^1 \leq T$. By Stieltjes integration by parts and Fubini's theorem,

$$\begin{aligned} R_t(T) &= \mathbb{E}\left(\int_t^T F(t,s)dN_s^1 \,\middle|\, \mathcal{F}_t\right) \\ &= F(t,T)\mathbb{P}(T^1 \leq T \,|\, \mathcal{F}_t) + r\int_t^T F(t,s)\mathbb{P}(T^1 \leq s \,|\, \mathcal{F}_t)ds. \end{aligned} \tag{30}$$

The pre-default value of the bond $B_t(T)$ is given by the value of the survival cash flow (28) plus the value of the recovery cash flow (29). If the survival probability (24) is analytically tractable then so is the corporate bond pricing problem.

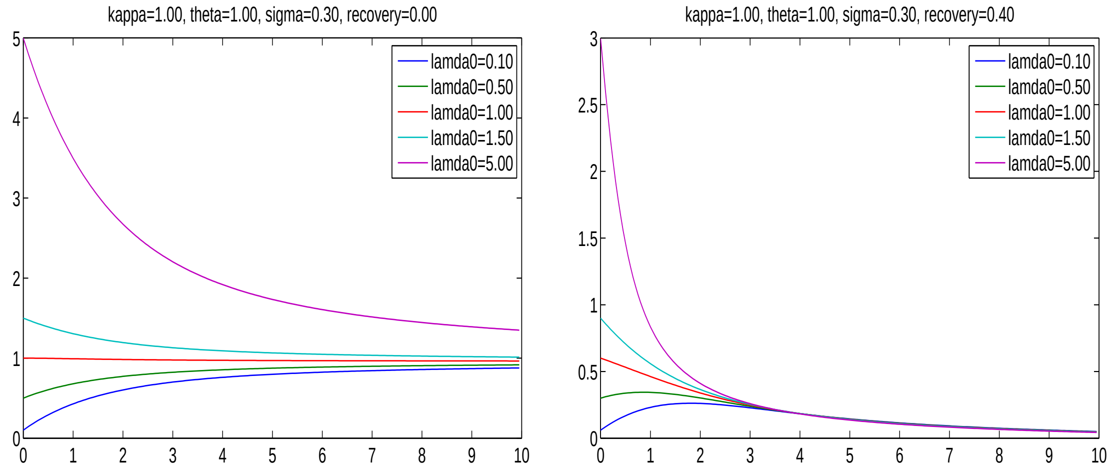

Figure 2: Term structure of zero coupon credit spreads (31) at $t = 0$ when N is an affine point process with intensity $\lambda = X$, where X follows a mean-reverting Feller diffusion process: $dX_t = \kappa(c - X_t)dt + \sigma\sqrt{X_t}dW_t$. That is, $K_0 = \kappa c$, $K_1 = -\kappa$, $H_0 = 0$, $H_1 = \sigma^2$ and $\delta = 0$ in the risk factor dynamics (12). The parameter κ controls the speed of mean reversion, c prescribes the mean reversion level, and σ governs the diffusive volatility of X. The condition $2\kappa c \geq \sigma^2$ guarantees positivity of X. The conditional survival probability $\mathbb{P}(T^1 > T \mid \mathcal{F}_t)$ is given by formula (25). The associated ODEs (26) and (27) can be solved analytically to give a closed formula for the bond price and credit spread. The risk-free rate $r = 5\%$, $\kappa = c = 1$ and $\sigma = 0.3$. *Left panel*: Zero recovery, i.e. $\ell = 1$. *Right panel*: Expected recovery 0.4, i.e. $\ell = 0.6$.

The zero coupon *credit spread* CS is the difference between the (promised) yield on a corporate zero bond and the yield on a government zero bond with matching face value and maturity. With $F(t, T)$ the price of the government bond,

$$CS_t(T) = -\frac{1}{T-t}\log\left(\frac{B_t(T)}{F(t,T)}\right) \tag{31}$$

at any time $t < T$ prior to default. Thus, the value $CS_t(T)$ is the excess yield at time t over the risk-free yield r demanded by the bond investor for bearing the default risk over the remaining term $T - t$. Figure 2 plots the term structure of credit spreads $T \to CS_0(T)$ when N is an affine point process driven by a state variable that follows a mean-reverting Feller diffusion process; see Feller (1951) and Section 3.2 above.

4.2 Corporate coupon bond

Most corporate bonds are issued with a coupon that stipulates a stream of interest payments. A corporate coupon bond with unit face value, annualized coupon rate c, coupon dates (t_m) and maturity T pays

- The coupon cC_m at each $t_m < T^1$ where C_m is the day count fraction for period m,

- The face value 1 at $T < T^1$,
- The recovery $(1-\ell^1)$ of face value at $T^1 \leq T$,
- The accrued coupon $\frac{T^1 - t_{m-1}}{\Delta_m} cC_m$ at T^1 if $t_{m-1} < T^1 \leq t_m$, where $\Delta_m = t_m - t_{m-1}$.

The coupon stream is a portfolio of zero-recovery zero bonds with maturities t_m and face values cC_m. Therefore, at any time t before default, the coupon stream has value

$$c \sum_{t_m \geq t} F(t, t_m) C_m \mathbb{P}(T^1 > t_m \,|\, \mathcal{F}_t).$$

The principal payment is a zero-recovery zero bond with maturity T and face value 1 and therefore its value is $F(t,T)\mathbb{P}(T^1 > T \,|\, \mathcal{F}_t)$ at any time t before default. The pre-default value of the recovery is $(1-\ell)R_t(T)$. Suppose the issuer defaults between two coupon dates. The accrued coupon covers the interest that has accumulated since the last coupon date. Typically, the accrued coupon is not paid separately, but may be subsumed into the recovery payment at default. Although artificial, here we treat it separately since it will play a role for credit swaps. The pre-default value of the accrued coupon is

$$c \sum_{t_m \geq t} \frac{C_m}{\Delta_m} \mathbb{E}\left(\int_{t \vee t_{m-1}}^{t_m} F(t,s)(s - t_{m-1}) dN_s^1 \,\middle|\, \mathcal{F}_t \right).$$

The value of the corporate coupon bond is given by the sum of the values of the four parts; at any time t before default we get

$$F(t,T)\mathbb{P}(T^1 > T \,|\, \mathcal{F}_t) + cV_t(T) + (1-\ell)R_t(T)$$

where $V_t(T)$ is the *risky DV01*, the value at any time t before default of a unit stream at coupon times (t_m) until $T^1 \wedge T$ plus any accruals, given by

$$V_t(T) = \sum_{t_m \geq t} \frac{C_m}{\Delta_m} \Big\{ F(t,t_m)\Delta_m - F(t, t \vee t_{m-1})(t \vee t_{m-1} - t_{m-1}) \mathbb{P}(T^1 \leq t \vee t_{m-1} \,|\, \mathcal{F}_t) - \int_{t \vee t_{m-1}}^{t_m} F(t,s)(1 - r(s - t_{m-1}))\mathbb{P}(T^1 \leq s \,|\, \mathcal{F}_t) ds \Big\}. \tag{32}$$

Now we are ready to analyze a credit swap.

4.3 Credit swap

A credit swap with unit notional, annualized swap spread S, premium payment dates (t_m) and maturity T is a bilateral contract in which

- The protection seller pays the default loss ℓ^1 at $T^1 \leq T$ (default leg),

- The protection buyer pays the swap spread SC_m at each $t_m < T^1$ plus any accruals (premium leg).

The value of the premium leg at any time t before default is

$$P_t(S) = SV_t(T) \tag{33}$$

where $V_t(T)$ is the risky DV01 given in formula (32). The value of the default leg at any time t before default is

$$D_t = \ell R_t(T) \tag{34}$$

where $R_t(T)$ is the value of a unit recovery payment at $T^1 \leq T$ given in formula (30). The fair spread S equates the values at inception of the default and premium legs. Since there is no cash flow at inception, the fair swap spread at inception date t is the solution $S = S_t(T)$ to the equation $D_t = P_t(S)$. Thus, $S_t(T) = \ell R_t(T)/V_t(T)$ at any time t prior to default. Note that, to derive this formula, we have only assumed that the firm's default time and recovery rate are independent, that the risk-free rate of interest is constant, and that the default risk of the protection seller is negligible.[4]

Consider an investor who buys protection at $t = 0$ for the period $[0, T]$ for a swap spread of $S_0(T)$. The mark-to-market value of the investor's position at time $t \geq 0$, when the market spread is $S_t(T)$, is given by $V_t(T)(S_t(T) - S_0(T)) = D_t - P_t(S_0(T))$ at any time t prior to default. The mark-to-market value of the protection seller's position is the negative of the buyer's value. Note that the mark-to-market value can be negative. This occurs when the credit quality of the reference name has improved since inception, and default protection is cheaper at current conditions.

Suppose we observe a firm's swap spreads in the credit swap market. We can extract the firm's (risk-neutral) survival probabilities from these spread quotes. That is, we fix a parametric family of intensity models, and choose model parameters so as to match the observed spreads as closely as possible. The calibrated model implies survival probabilities for all maturities. This procedure is illustrated in Figure 3, where we extract survival probabilities for the investment bank Lehman Brothers and the telecommunications firm Comcast, based on market spreads for several maturities on September 2, 2008. Lehman's spread term structure is decreasing, indicating the market's awareness of Lehman's difficulties at that time. Lehman defaulted on September 15, 2008.

The market implied survival probabilities have multiple applications. For example, they can be used to design credit trading strategies. One strategy is to exploit different valuations of a firm's creditworthiness that may exist in different markets. For instance, the swap market implied survival probabilities can be contrasted with survival probabilities extracted from the firm's equity prices, in order to identify those firms for which the gap is

[4]Counterparty risk, the risk that the protection seller fails with the reference entity, has become a serious issue in the current market environment. Intuitively, it lowers the fair swap spread $S_t(T)$ calculated under the assumption that counterparty risk is negligible.

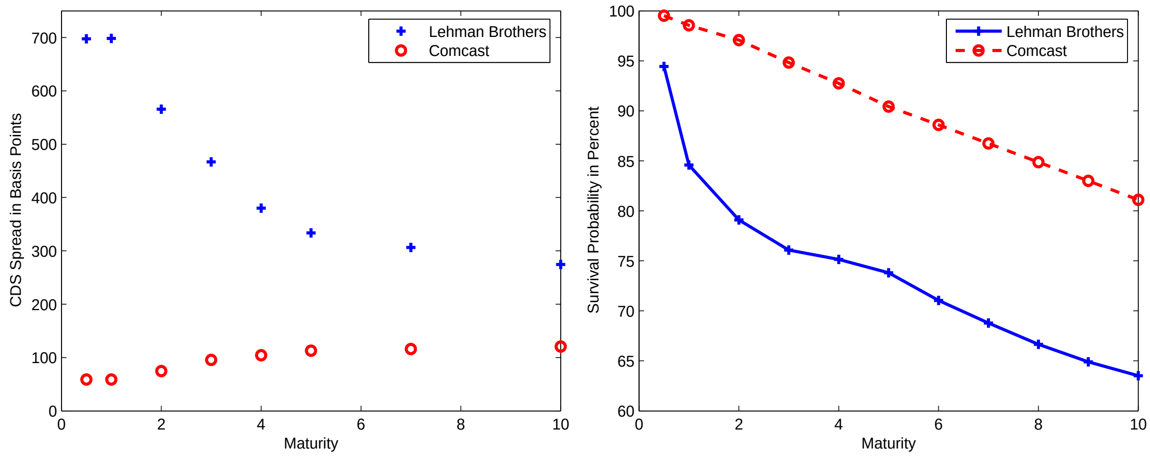

Figure 3: Extracting the survival probabilities of Lehman Brothers and Comcast from market credit swap spreads. *Left panel:* Mid-market spread quotes on September 2, 2008, for each of several maturities. Source: Barclays Capital. To fit these quotes perfectly, we assume that the process N has piece-wise constant intensity $\lambda(t)$, whose intervals of constancy are given by the inter-maturity times of the given spread quotes. With a deterministic intensity, N is a non-homogeneous Poisson process and the survival probability $\mathbb{P}(T^1 > T \mid \mathcal{F}_t) = \exp(-\int_t^T \lambda(s)ds)$. We sequentially determine the values of the intensity on its intervals of constancy from spread quotes of the corresponding maturities, starting with the shortest maturity, by finding the intensity value that equates a quote to the model spread for a fixed expected loss $\ell = 0.6$, the standard industry value. This inverse problem is usually well-posed. If it is not, then that indicates arbitrage opportunities. *Right panel:* Survival probabilities $\mathbb{P}(T^1 > T)$ implied by the calibrated model.

relatively large. The investor can bet on the gap to narrow over time. Another important application of the implied survival probabilities is the valuation of other credit-sensitive securities that are issued by or referenced on the firm in question. One such instrument is considered in the next section.

4.4 Forward credit swap

In a regular (spot) credit swap the parties agree on a price for protection that starts immediately. In a forward credit swap, the parties agree on a price for protection that starts at a specified time in the future. More precisely, in a forward credit swap with unit notional, annualized forward spread S, premium payment dates (t_m), forward start date U and maturity date $T > U$,

- The protection seller pays the default loss ℓ^1 at $T^1 \in [U, T]$ (default leg),
- The protection buyer pays the forward swap spread SC_m at each $t_m \in (U, T^1)$ plus any accruals (premium leg).

Note that cash flows occur only between times U and T. Of course, if U is the time at which the spread is negotiated, then the forward swap is a spot swap. Also note that the contract is canceled ("knocked out") in case the reference issuer defaults before U.

At any time t before default, the premium leg has value

$$P_t(S) = S(V_t(T) - V_t(U))$$

The default leg covers the loss over $[U, T]$ and is equal to the difference between the default leg of a spot swap maturing at T and the default leg of a spot swap maturing at U. Therefore, its pre-default value must be equal to the difference of the pre-default values of the corresponding premium legs:

$$D_t = S_t(T)V_t(T) - S_t(U)V_t(U)$$

The fair forward spread at time t is the solution $S = S_t(U, T)$ to the equation $D_t = P_t(S)$ so at any time t prior to default

$$S_t(U, T) = \frac{S_t(T)V_t(T) - S_t(U)V_t(U)}{V_t(T) - V_t(U)}.$$

Using this formula, we can estimate the forward spread $S_t(U, T)$ from the market spot swap spreads $S_t(U)$ and $S_t(T)$, and the corresponding risky DV01s $V_t(U)$ and $V_t(T)$, which are calculated via formula (32) based on the market implied survival probabilities.

4.5 Credit swaption

A forward swap obligates the parties to enter into the swap at the forward start date. An option on a credit swap, or swaption with maturity T imparts the right but not the obligation to enter into a swap at a future expiration date $U < T$ at a strike spread S. A payer swaption gives the right to buy protection at U for a spread S. As a forward swap, it can include a knock-out provision, which would stipulate that the contract is canceled in case the reference issuer defaults between the trade date and option's expiry date U. While a payer swaption allows the investor to monetize widening spreads (deteriorating credit quality), a receiver swaption imparts the right to sell protection at a spread S, and thus allows the investor to bet on tightening spreads.

5 Multi-name credit derivatives

Corporate defaults are correlated because firms are exposed to common or correlated economic risk factors, see Duffie, Saita & Wang (2006). The movements of the risk factors generate correlated changes in firms' conditional default rates. This relationship is illustrated in Figure 4, which shows the number of defaults of Moody's rated firms in any given year between January 1970 and September 2008. Some of the risk factors may not

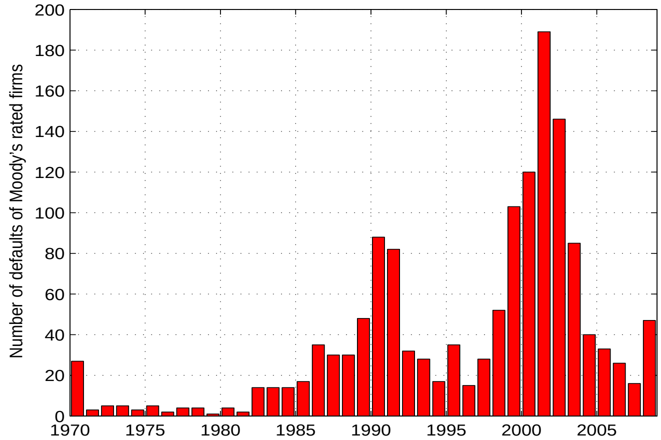

Figure 4: Annual defaults of Moody's rated firms between January 1970 and September 8, 2008. Source: Moody's Default Risk Service. The peak in 1970 represents a cluster of 24 railway defaults triggered by the collapse of Penn Central Railway on June 21, 1970. The fallout of the 1987 crash is indicated by the peak in the early nineties. The burst of the internet bubble caused many defaults during 2001-2002. From a trough in 2007, default rates started to increase significantly in 2008.

be observable, and the uncertainty associated with unobservable factors generates an additional channel of default correlation, see Collin-Dufresne, Goldstein & Helwege (2009), Giesecke (2004), Duffie, Eckner, Horel & Saita (2009). Furthermore, a default may have a contagious impact on the surviving firms that is channeled through the complex web of contractual relationships in the economy, see Azizpour, Giesecke & Schwenkler (2009). The idea is that the failure of a firm tends to weaken the others. The ongoing crisis in credit markets provides evidence of the existence of contagion phenomena.

Multi-name or portfolio credit derivatives are referenced on a pool of 5 to several hundred firms. They are often structured as swaps, very much like a credit swap referenced on a single obligation, or single-name swap. The instruments are designed to facilitate the transfer of the correlated default risk in the reference portfolio. From a modeling perspective, the challenge is to capture the sources of default clustering mentioned above while maintaining the computational tractability of the valuation relations, which require the distributions of the point processes describing portfolio defaults and losses.

There are two different modeling philosophies. In a *bottom-up approach*, the default processes of the portfolio constituents are the modeling primitives, see Duffie & Garleanu (2001), Eckner (2009), Mortensen (2006), Papageorgiou & Sircar (2007) and others. That is, for each constituent i we specify a counting process, $N(i)$ say, most conveniently through an intensity $\lambda(i)$. The dependence between the constituent intensity processes reflects the default correlation in the portfolio. The first jump of $N(i)$ models the default time of firm i, as in Section 4 above. Single-firm default probabilities are given by a formula analogous to (24) in terms of $\lambda(i)$. The distribution of the portfolio default counting process $N = \sum_i (N(i) \wedge 1)$ takes a more complex form; see Giesecke & Zhu (2010) for a general formula.

In an alternative *top-down approach*, the portfolio default counting process N is specified in terms of an intensity λ, without reference to the portfolio constituents. Examples of this approach are Arnsdorf & Halperin (2008), Cont & Minca (2008), Ding, Giesecke & Tomecek (2009), Errais et al. (2010), and others. The default correlation among the portfolio constituents is reflected in the dynamics of λ. For example, we could specify N as an affine point process; see Section 3.2 above. The jumps in the intensity would reflect the contagious impact of an event, while the diffusive movements of the inter-arrival intensity would reflect firms' exposure to a risk factor. The inversion formula (3) would lead to the distribution of the portfolio default process. Further steps are required to obtain the survival probabilities of the portfolio constituents. Giesecke, Goldberg & Ding (2011) propose random thinning to disintegrate λ into the constituent intensities $\lambda(i)$, and develop a corresponding formula for the constituent survival probabilities. This formula leads to a scheme for estimating constituent hedge sensitivities for multi-name credit derivatives. These sensitivities are important in practice.

Below we take a top-down approach, with N representing the portfolio default process, to analyze index and tranche swaps. These swaps are referenced on a portfolio of single-name credit swaps rather than bonds or loans. The collection of tranche swaps referenced on a given portfolio is sometimes called a *collateralized debt obligation*, or CDO.

5.1 Index swaps

An index swap is the most liquid portfolio derivative. It is based on a portfolio of n single-name credit swaps. The constituent swaps have common notional that we normalize to 1, common maturity date T and common premium payment dates (t_m). A default in the reference obligation of a constituent swap translates into a loss for the constituent swap protection seller. This loss is recorded by the portfolio loss process $L = \sum_{k=0}^{N} \ell^k$, where $\ell^k \in [0, 1]$ is the normalized random loss at the kth default. In an index swap with swap spread S, the cash flows are as follows:

- The protection seller covers portfolio losses as they occur, i.e. the increments of the portfolio loss process L (default leg),
- The protection buyer pays $SC_m(n - N_{t_m})$ at each date t_m (premium leg).

The protection seller assumes exposure to the correlated default risk associated with the reference portfolio. There are multiple families of standard reference portfolios, called indices. The CDX family covers North American issuers. There are various CDX subindices, including the Investment Grade (125 constituents) and High Yield (100 constituents) indices. The iTraxx index family represents European and Asian portfolios.

The standardization of reference portfolios helps to increase liquidity in the index swap market, which benefits all market participants.

The value D_t at time $t \le T$ of the index swap default leg is given by the discounted cumulative losses. Note that, unlike a single-name swap, an index swap is not terminated at a default event. By integration by parts and Fubini, we get

$$\begin{aligned} D_t &= \mathbb{E}\left(\int_t^T F(t,s)dL_s \,\middle|\, \mathcal{F}_t\right) \\ &= F(t,T)\mathbb{E}(L_T \,|\, \mathcal{F}_t) - L_t + r\int_t^T F(t,s)\mathbb{E}(L_s \,|\, \mathcal{F}_t)\, ds. \end{aligned} \tag{35}$$

The value at time $t \le T$ of the premium leg is given by

$$P_t(S) = S \sum_{t_m \ge t} F(t,t_m) C_m (n - \mathbb{E}(N_{t_m} \,|\, \mathcal{F}_t)). \tag{36}$$

The fair index swap spread at time t is the solution $S = S_t(T)$ to the equation $D_t = P_t(S)$. Formulae (35) and (36) indicate that the index spread depends only on expected defaults and losses for horizons between t and T. Figure 5 plots the term structure $T \to S_0(T)$ of index swap spreads when the default process N is a self-exciting affine point process.

5.2 Tranche swaps

The index protection seller is exposed to the correlated default risk associated with the entire portfolio. Investors seeking narrower risk profiles can trade a tranche swap referenced on the portfolio. A tranche swap is specified by a lower attachment point $\underline{K} \in [0,1]$ and an upper attachment point $\overline{K} \in (\underline{K}, 1]$. The product of the difference $K = \overline{K} - \underline{K}$ and the portfolio notional n is the tranche notional. In a tranche swap with upfront rate G and swap spread S, the cash flows are as follows:

- The protection seller covers tranche losses as they occur, i.e. the increments of the tranche loss $U_t = (L_t - \underline{K}n)^+ - (L_t - \overline{K}n)^+$ (default leg),
- The protection buyer pays GKn at inception and $SC_m(Kn - U_{t_m})$ at each date t_m (premium leg, assuming $\overline{K} < 1$).

Figure 6 shows a sample path of the loss process and the corresponding payments of the protection seller. Unlike an index protection seller, a tranche protection seller is exposed only to a slice of the loss that is specified by the lower and upper attachment points. The higher the lower attachment point, the lower is the risk of payouts for the protection seller, and the lower is the fee that the protection seller can require. This is because the losses must first eat through the subordinated tranches, which provide a safety cushion. The equity tranche, whose lower attachment point is zero, has no safety cushion: the equity protection seller must bear any losses until the equity notional is exhausted. It

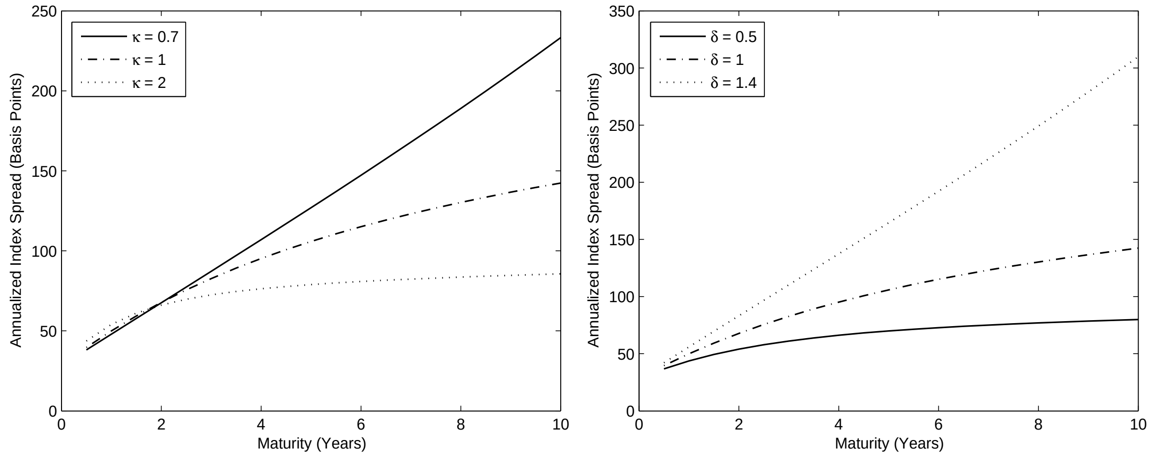

Figure 5: Term structure of index swap spreads at $t = 0$ when N is a self-exciting affine point process with intensity $\lambda = X$, where X follows the process $dX_t = \kappa(c - X_t)dt + \delta dJ_t$, see Errais et al. (2006). That is, $K_0 = \kappa c$, $K_1 = -\kappa$ and $H_0 = H_1 = 0$ in the risk factor dynamics (12), and N is a Hawkes process. The intensity jumps in response to a default, with random jump sizes that are drawn independently of one another from a uniform distribution on $\{0.4, 0.6, 0.8, 1\}$. The Laplace transform $\varphi(u, t, T)$ of N is given by formula (20). The corresponding satisfy the ODEs (19) and (18) are solved numerically. The mean $\mathbb{E}(N_T \,|\, \mathcal{F}_t) = N_t - \partial_u \varphi(u, t, T)|_{u=0}$. The mean loss $\mathbb{E}(L_T \,|\, \mathcal{F}_t) = \ell \mathbb{E}(N_T \,|\, \mathcal{F}_t)$, assuming that the loss variables ℓ^k are drawn independently of one another and independently of N and J from a common distribution with mean ℓ. We take that distribution to be the uniform distribution on $\{0.4, 0.6, 0.8, 1\}$. The number of names $n = 125$. The risk-free rate $r = 5\%$. Unless stated otherwise, we set $X_0 = 0.5$, $c = \delta = \kappa = 1$. *Left panel*: Index spread term structure for each of several values of κ, which governs the exponential decay of the intensity after an event. The higher κ, the faster the intensity reverts back to the base intensity c after a default. That is, the quicker the negative impact of a default on the economy fades away, and the lower the risk of a default cluster. Thus, index spreads are decreasing in κ. *Right panel*: Index spread term structure for each of several values of δ, which governs the sensitivity of the intensity to a default. The higher δ, the larger the mean impact of a default on the economic state, and the higher the probability of a default cluster, which generates large losses. Thus, index spreads are increasing in δ.

is the riskiest tranche, and commands the highest premium. In a customized or bespoke deal, the attachment points are set so that an investor's target risk profile is met.

The value at time t of the default leg is

$$D_t(\underline{K}, \overline{K}) = F(t, T)\mathbb{E}(U_T \,|\, \mathcal{F}_t) - U_t + r \int_t^T F(t, s)\mathbb{E}(U_s \,|\, \mathcal{F}_t)\, ds. \tag{37}$$

This formula is analogous to formula (35) for the value of an index swap default leg. The latter can be viewed as the default leg of a tranche swap for which $\underline{K} = 0$ and $\overline{K} = 1$.

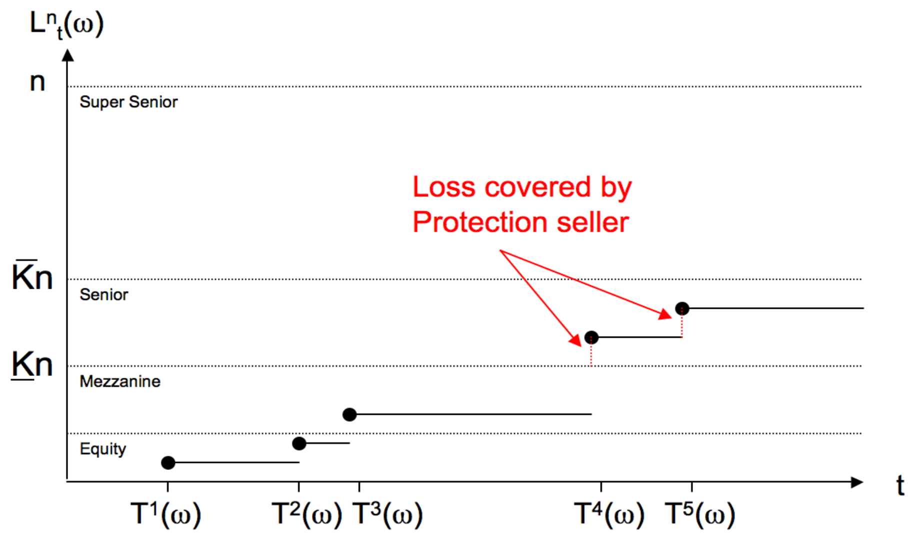

Figure 6: Paths of the portfolio loss process L and the ordered default times (T^k), and payments of the protection seller in a senior tranche swap. The protection seller is exposed to the correlated default risk in the reference portfolio. The risk profile is determined by the location of the attachment points $\underline{K}$ and $\overline{K}$. The riskiest tranche is the equity tranche, for which $\underline{K} = 0$ and $\overline{K}$ is typically 3 or 10%. The least risky tranche is the super senior tranche, for which $\overline{K} = 1$ and $\underline{K}$ is typically 30 or 35%. This tranche bears losses only if all the other tranches are wiped out. The collection of all tranches is called a collateralized debt obligation, or CDO.

The value of the premium leg is given by

$$P_t(\underline{K}, \overline{K}, G, S) = GKn + S \sum_{t_m \geq t} F(t, t_m) C_m (Kn - \mathbb{E}(U_{t_m} \,|\, \mathcal{F}_t)). \tag{38}$$

For a fixed upfront payment rate G, the fair tranche spread S is the solution $S = S_t(\underline{K}, \overline{K}, G, T)$ to the equation $D_t(\underline{K}, \overline{K}) = P_t(\underline{K}, \overline{K}, G, S)$. Similarly, for a fixed tranche spread S, the fair tranche upfront rate G is the solution $G = G_t(\underline{K}, \overline{K}, S, T)$ to the equation $D_t(\underline{K}, \overline{K}) = P_t(\underline{K}, \overline{K}, G, S)$. The fair spread and upfront rate depend only on the value of *call spreads*

$$F(t, s)\mathbb{E}(U_s \,|\, \mathcal{F}_t) = F(t, s)(\mathbb{E}((L_s - \underline{K}n)^+ \,|\, \mathcal{F}_t) - \mathbb{E}((L_s - \overline{K}n)^+ \,|\, \mathcal{F}_t))$$

on the portfolio loss L_s with strikes $\underline{K}n$ and $\overline{K}n$ and maturities s between t and T. Figure 7 shows the term structure $T \to S_0(3\%, 7\%, 0, T)$ of mezzanine tranche spreads when the default process N is a self-exciting affine point process, as in Figure 5.

In analogy to the single-name swap case, we can calibrate a parametric model of the risk-neutral intensity λ of the portfolio default process N from market spreads of index and tranche swaps with different maturities and attachment points, all referenced on the same portfolio. The calibration procedure is illustrated in Figure 8 for the self-exciting

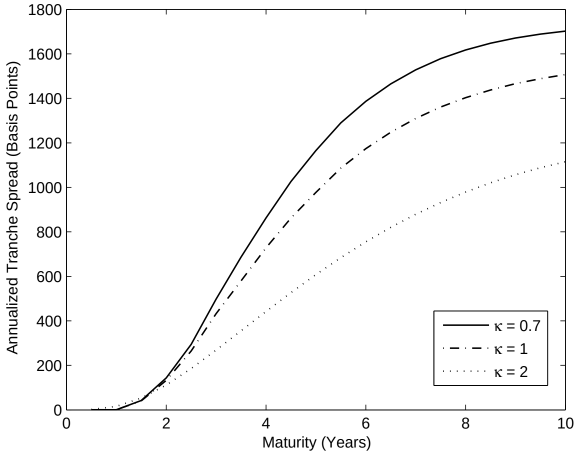

Figure 7: Term structure of mezzanine tranche spreads $S_0(3\%, 7\%, 0, T)$ at $t = 0$ when N is a self-exciting affine point process with intensity $\lambda = X$, where X follows the process $dX_t = \kappa(c - X_t)dt + \delta dJ_t$, see Errais et al. (2006). That is, $K_0 = \kappa c$, $K_1 = -\kappa$ and $H_0 = H_1 = 0$ in the risk factor dynamics (12), and N is a Hawkes process, as in Figure 5. The distribution of $N_T - N_t$ is obtained via formula (3), and used to calculate the Laplace transform $\mathbb{E}(e^{-u(L_T - L_t)} \mid \mathcal{F}_t) = \sum_{k \geq 0} (f(u))^k \mathbb{P}(N_T - N_t = k \mid \mathcal{F}_t)$ of the loss, assuming that the loss variables ℓ^k are drawn independently of one another and independently of N and J from a common distribution with Laplace transform $f(u)$. We take that distribution to be the uniform distribution on $\{0.4, 0.6, 0.8, 1\}$. The Laplace transform of the loss is inverted to obtain the loss distribution, which is used to calculate the (undiscounted) call value $\mathbb{E}((L_s - c)^+ \mid \mathcal{F}_t) = \int_c^n (x - c) d\mathbb{P}(L_s \leq x \mid \mathcal{F}_t)$ for $0 \leq c \leq n$. The number of names $n = 125$. The risk-free rate $r = 5\%$. Unless stated otherwise, we set $X_0 = 0.5$, $c = \delta = \kappa = 1$. *Left panel*: Mezzanine spread term structure for each of several values of κ, which governs the exponential decay of the intensity after an event. The higher κ, the faster the intensity reverts back to the base intensity c after a default. *Right panel*: Mezzanine spread term structure for each of several values of δ, which governs the propensity of defaults to cluster. The more senior the tranche (i.e., the higher the lower attachment point $\underline{K}$), the more sensitive is the tranche spread to δ. This is because the mass in the tail of the loss distribution is increasing in δ, and the more senior the tranche the more sensitive it is to scenarios with large losses.

affine point process considered in Figures 5 and 7, for which λ is an affine function of an affine jump diffusion. The calibrated intensity model induces the market-implied risk-neutral distributions of the portfolio default and loss processes. It is used to estimate the arbitrage-free value of a non-traded portfolio derivative, for example a tranche referenced on the given portfolio but with non-standard attachment points. There is often model risk associated with this procedure, in that there may be many distinct parametric intensity models that fit the given market data but imply different prices for a given non-traded portfolio derivative.

Contract	MaBid	MaAsk	Model
0-10	70.50%	70.75%	71.48%
10-15	34.25%	34.50%	32.74%
15-25	316.00	319.00	311.43
25-35	79.00	81.00	77.34
Index	262.85	263.10	262.97
AAPE			2.24%

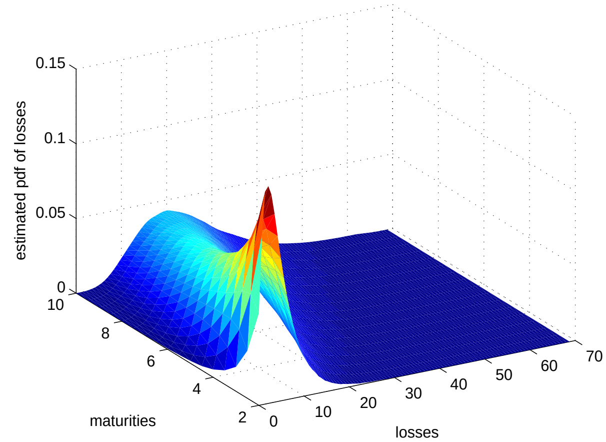

Figure 8: Index and tranche market calibration at $t = 0$ when N is a self-exciting affine point process with intensity $\lambda = X$, where X follows the process $dX_t = \kappa(c - X_t)dt + \delta dJ_t$, as analyzed in Figures 5 and 7. The loss variables ℓ^k are drawn independently of one another and independently of N and J from a common uniform distribution over $\{a, b\}$ with $0 < a < b < 1$. Consistent with market practice, we set the expected loss at default $\ell = \mathbb{E}(\ell^k) = 0.6$. The parameter vector to be calibrated is $\theta = (X_0, c, \kappa, \delta, a)$. We solve the optimization problem $\min_{\theta \in \Theta} \sum_i (\text{MaMid}(i) - \text{Model}(i, \theta))^2 / (\text{MaAsk}(i) - \text{MaBid}(i))^2$, where $\Theta = [0, 5]^4 \times [0.2, 0.6]$ and the sum ranges over the data points. The market mid quote MaMid is the arithmetic average of the observed MaAsk and MaBid quotes. The risk-free rate $r = 5\%$. More details are given in Giesecke & Kim (2007). *Left panel*: Market bid and ask quotes for the 5 year maturity index and tranches on the CDX High Yield portfolio of $n = 100$ names, observed on 5/11/2007, together with model calibrated values at the optimal parameter vector $\theta^* = (0.75, 1.60, 2.58, 2.94, 0.24)$. The quotes are in basis points except for the $[0, 10]$ and $[10, 15]$ percent tranches, which are quoted in terms of a percentage upfront fee. We report the average absolute percentage fitting error AAPE. *Right panel*: Smoothed portfolio loss distribution $\mathbb{P}(L_T \in dx)$ for T between 2 and 10 years implied by the calibrated model.

5.3 Other instruments

There is a range of other portfolio credit derivatives that parallel the instruments in the single-name space. In particular, index and tranche swaps are also traded on a forward basis. The forward spread curve can be constructed from the spot spread curve using the arbitrage argument given in Section 4.4 for the single-name case. There are also options on index and tranche swaps, which are analyzed in Ding et al. (2009).

A collateralized debt obligation can be structured in several ways. Above we have analyzed a *synthetic* CDO, which is backed by credit swaps. In a *cash* CDO, the pool of reference securities contains straight bonds, loans or other credit obligations. While the rationale of the cash CDO is similar to that of the synthetic CDO, the cash structure often has additional features that complicate the analysis. For example, cash instruments such as bonds or loans generate interest payments that increase collateral. Mortgages can be retired before their maturity, generating additional pre-payment risk. In order to account for such features, simulation methods as in Duffie & Garleanu (2001) and Giesecke & Kim (2011) are required.

5.4 Risk analysis

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 index and tranche swaps. The risk analysis problem is distinct from the valuation problem discussed above: it 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. Giesecke & Kim (2011) develop, analyze and test an acceptance/rejection resampling scheme to estimate this distribution from historical default timing data.

References

Arnsdorf, Matthias & Igor Halperin (2008), 'BSLP: markovian bivariate spread-loss model for portfolio credit derivatives', *Journal of Computational Finance* **12**, 77–100.

Azizpour, Shahriar, Kay Giesecke & Baeho Kim (2011), 'Premia for correlated default risk', *Journal of Economic Dynamics and Control* **35**(8), 1340–1357.

Azizpour, Shahriar, Kay Giesecke & Gustavo Schwenkler (2009), Exploring the sources of default clustering. Working Paper, Stanford University.

Berndt, Antje, Rohan Douglas, Darrell Duffie, Mark Ferguson & David Schranz (2005), Measuring default risk premia from default swap rates and EDF's. Working Paper, Stanford University.

Beskos, Alexander & Gareth Roberts (2005), 'Exact simulation of diffusions', *Annals of Applied Probability* **15**(4), 2422–2444.

Brémaud, Pierre (1980), *Point Processes and Queues – Martingale Dynamics*, Springer-Verlag, New York.

Carr, Peter & Liuren Wu (2004), 'Time-changed Lévy processes and option pricing', *Journal of Financial Economics* **71**, 113–141.

Chen, Nan (2009), Localization and exact simulation of brownian motion driven stochastic differential equations. Working Paper, Chinese University of Hong Kong.

Collin-Dufresne, Pierre, Robert Goldstein & Jean Helwege (2009), How large can jump-to-default risk premia be? modeling contagion via the updating of beliefs. Working Paper, Columbia University.

Cont, Rama & Andreea Minca (2008), Extracting portfolio default rates from CDO spreads. Working Paper, Columbia University.

Cox, John, Jonathan Ingersoll & Stephen Ross (1985), 'A theory of the term structure of interest rates', *Econometrica* **53**, 385–408.

Dellacherie, Claude & Paul-Andre Meyer (1982), *Probabilities and Potential*, North Holland, Amsterdam.

Ding, Xiaowei, Kay Giesecke & Pascal Tomecek (2009), 'Time-changed birth processes and multi-name credit derivatives', *Operations Research* **57**(4), 990–1005.

Duffie, Darrell (1996), *Dynamic Asset Pricing Theory*, second edn, Princeton University Press, Princeton, New Jersey.

Duffie, Darrell, Andreas Eckner, Guillaume Horel & Leandro Saita (2009), 'Frailty correlated default', *Journal of Finance* **64**, 2089–2123.

Duffie, Darrell, Damir Filipovic & Walter Schachermayer (2003), 'Affine processes and applications in finance', *Annals of Applied Probability* **13**, 984–1053.

Duffie, Darrell, Jun Pan & Kenneth Singleton (2000), 'Transform analysis and asset pricing for affine jump-diffusions', *Econometrica* **68**, 1343–1376.

Duffie, Darrell, Leandro Saita & Ke Wang (2006), 'Multi-period corporate default prediction with stochastic covariates', *Journal of Financial Economics* **83**(3), 635–665.

Duffie, Darrell & Nicolae Garleanu (2001), 'Risk and valuation of collateralized debt obligations', *Financial Analysts Journal* **57**(1), 41–59.

Eckner, Andreas (2007), Risk premia in structured credit derivatives. Working Paper, Stanford University.

Eckner, Andreas (2009), 'Computational techniques for basic affine models of portfolio credit risk', *Journal of Computational Finance* **15**, 63–97.

Errais, Eymen, Kay Giesecke & Lisa Goldberg (2010), 'Affine point processes and portfolio credit risk', *SIAM Journal on Financial Mathematics* **1**, 642–665.

Feller, William (1951), 'Two singular diffusion problems', *Annals of Mathematics* **54**, 173–182.

Föllmer, Hans & Alexander Schied (2004), *Stochastic Finance: An Introduction in Discrete Time*, De Gruyter, Berlin.

Giesecke, Kay (2004), 'Correlated default with incomplete information', *Journal of Banking and Finance* **28**, 1521–1545.

Giesecke, Kay & Baeho Kim (2007), Estimating tranche spreads by loss process simulation, *in* S. G.Henderson, B.Biller, M.-H.Hsieh, J.Shortle, J. D.Tew & R. R.Barton, eds, 'Proceedings of the 2007 Winter Simulation Conference', IEEE Press, pp. 967–975.

Giesecke, Kay & Baeho Kim (2011), 'Risk analysis of collateralized debt obligations', *Operations Research* **59**(1), 32–49.

Giesecke, Kay, Baeho Kim & Shilin Zhu (2010), Monte Carlo algorithms for default timing problems. *Management Science*, forthcoming.

Giesecke, Kay & Dmitry Smelov (2010), Exact sampling of jump-diffusions. Working Paper, Stanford University.

Giesecke, Kay, Hossein Kakavand & Mohammad Mousavi (2010), Exact simulation of point processes with stochastic intensities. *Operations Research*, forthcoming.

Giesecke, Kay, Hossein Kakavand, Mohammad Mousavi & Hideyuki Takada (2010), 'Exact and efficient simulation of correlated defaults', *SIAM Journal on Financial Mathematics* **1**, 868–896.

Giesecke, Kay, Lisa Goldberg & Xiaowei Ding (2011), 'A top-down approach to multi-name credit', *Operations Research* **59**(2), 283–300.

Giesecke, Kay & Shilin Zhu (2010), Transform analysis for point processes and applications in credit risk. *Mathematical Finance*, forthcoming.

Glasserman, Paul & Nicolas Merener (2003), 'Numerical solution of jump-diffusion LIBOR market models', *Finance and Stochastics* **7**, 1–27.

Karatzas, Ioannis & Steven E. Shreve (1988), *Brownian Motion and Stochastic Calculus*, Graduate Texts in Mathemathics; 113, Springer-Verlag, New York.

Leippold, Markus & Liuren Wu (2002), 'Asset pricing under the quadratic class', *Journal of Financial and Quantitative Analysis* **37**, 271–295.

Lewis, Peter & George Shedler (1979), 'Simulation of nonhomogeneous Poisson processes by thinning', *Naval Logistics Quarterly* **26**, 403–413.

Meyer, Paul-André (1971), Démonstration simplifée d'un théorème de Knight, *in* 'Séminaire de Probabilités V, Lecture Note in Mathematics 191', Springer-Verlag Berlin, pp. 191–195.

Mortensen, Allan (2006), 'Semi-analytical valuation of basket credit derivatives in intensity-based models', *Journal of Derivatives* **13**, 8–26.

Ogata, Yosihiko (1981), 'On Lewis' simulation method for point processes', *IEEE Transactions on Information Theory* **27**, 23–31.

Papageorgiou, Evan & Ronnie Sircar (2007), 'Multiscale intensity models and name grouping for valuation of multi-name credit derivatives', *Applied Mathematical Finance* **15**(1), 73–105.

Protter, Philip (2004), *Stochastic Integration and Differential Equations*, Springer-Verlag, New York.

Sarig, Oded & Arthur Warga (1989), 'Some empirical estimates of the risk structure of interest rates', *Journal of Finance* **44**, 1351–1360.

Staum, Jeremy (2008), Incomplete markets, *in* J.Birge & V.Linetsky, eds, 'Handbooks in Operations Research and Management Science', Vol. 15, Elsevier, chapter 12.