A TOP-DOWN APPROACH TO MULTI-NAME CREDIT

Kay Giesecke* Lisa R. Goldberg†
Stanford University *MSCI Barra*

February 16, 2005; this draft June 6, 2008‡

Abstract

A multi-name credit derivative is a security tied to an underlying portfolio of corporate bonds or other credit-sensitive securities. It enables investors to buy and sell protection against the default losses in the portfolio. The value of a multi-name derivative depends on the distribution of portfolio loss at multiple horizons. Intensity-based models of the loss point process that are specified without reference to the portfolio constituents determine this distribution in terms of few economically meaningful parameters, and lead to tractable credit derivatives valuation relations that can be addressed by a variety of efficient methods. This paper proposes *random thinning* to extend the reach of these models beyond the portfolio level. Random thinning decomposes the portfolio loss process into the sum of its constituent loss processes, and allocates aggregate portfolio risk to sub-portfolios. We show that any loss process can be thinned, and that the associated thinning process is a probabilistic model for the next-to-default. We derive a formula for the constituent default probability in terms of the thinning process and the portfolio intensity, and show how to evaluate it for a large family of portfolio intensity models. This formula facilitates consistent pricing and calibration of constituent and portfolio credit derivatives, which we demonstrate by fitting a familiar stand-alone model of the portfolio loss to market rates of CDX index, tranche and constituent single-name credit swaps.

---

*Department of Management Science & Engineering, Stanford University, Stanford, CA 94305-4026, Phone (650) 723 9265, Fax (650) 723 1614, email: giesecke@stanford.edu, web: www.stanford.edu/∼giesecke.

†MSCI Barra, 2100 Milvia Street, Berkeley, CA 94704-1113, Phone (510) 649 4601, Fax (510) 848 0954, email: lisa.goldberg@mscibarra.com.

‡We thank Greg Anderson, Aaron Brown, Eymen Errais, Steve Evans, Igor Halperin, Steven Hutt, Jyh-Huei Lee, Thorsten Schmidt and Pascal Tomecek for illuminating discussions, and seminar participants at the 2007 Applied Probability Society Conference, the 2007 CreditMinds Conference, the 2007 Decision and Risk Analysis Conference, Fields Institute for Research in Mathematical Sciences, the 2007 German Mathematical Society Annual Conference, the 2006 ICBI Global Derivatives and Trading Conference, the 2006 and 2007 INFORMS Annual Meetings, J.P. Morgan, the 2007 Lévy models in Credit Conference, Morgan Stanley, Natixis, the 2007 Newton Institute Conference on Developments in Quantitative Finance, the 2006 Numerical Methods for Finance Conference, the 2006 Quantitative Methods in Finance Conference, the 2006 Quant Congress and the 2006 Stanford-Tsukuba Workshop on Financial Engineering for comments. We are deeply grateful to Xiaowei Ding for his contributions to the development of the material in Section 5 and to Stefan Weber for his contributions to the material in Appendix A.

1 Introduction

A multi-name or portfolio credit derivative is a financial instrument tied to a portfolio of defaultable securities, such as loans, corporate bonds or credit swaps referenced on these bonds. Its payoff is a function of the aggregate default loss in the underlying portfolio. For example, an index swap pays any losses before the contract maturity, while a tranche swap covers the losses that exceed a fixed threshold, up to a maximum amount. Portfolio credit derivatives play an important role in the financial market since they allow credit investors such as banks or asset managers to trade default insurance on the reference portfolio. Yet their analysis is challenging due to the complex economic phenomena that underpin correlated corporate default risk.

One approach is to view a multi-name instrument as a derivative written on the point process that models portfolio loss. Reduced-form or intensity-based models of the loss process that are specified without reference to the constituents have received significant interest for several reasons. First, they naturally extend the reduced-form models that are widely used to analyze corporate bonds and other single-name derivatives. Second, they allow the researcher to describe salient empirical phenomena including cyclical default behavior, contagion and frailty in terms of a concise set of economically meaningful parameters. Third, they lead to valuation formulae for standard and exotic portfolio derivatives that are computationally tractable, and that are largely independent of the number of assets in the reference portfolio. The formulation of the pricing problem supports the application of a range of computational techniques, including semi-analytical, transform, simulation, tree and PDE methods. Finally, they can accurately fit daily market index and tranche swap rates, as well as time series of swap rates.

There is an important catch, however. Since it is specified without reference to the constituent firms, a stand-alone reduced-form model of the portfolio loss process is silent about the component risks. Therefore, it cannot be used to address applications involving the constituent securities, for example single-name market calibration or constituent risk analysis and hedging. In this article, we show how to complement a stand-alone reduced-form model of the portfolio loss process to facilitate these important applications. We propose *random thinning* to decompose the default counting process driving portfolio loss into the sum of its constituent processes. Random thinning allocates to a sub-portfolio a fraction of the portfolio intensity, which is the primitive in a stand-alone model of portfolio loss. It generates sub-portfolio and single-name intensities that reflect the name dependence structure specified on the portfolio level. The resulting sub-models capture, for example, the cyclical, contagion and frailty effects in the ambient portfolio. They also treat sub-portfolio intersections consistently. This feature is important for the analysis of overlapping portfolios, which often occur in practice.

We show that any stand-alone model of the portfolio loss process can be thinned. The random thinning process that allocates the portfolio intensity to the constituent names has a natural interpretation. We demonstrate that it is a model for the next-to-default: its value represents the conditional probability that a firm defaults next, given default is imminent. We derive a formula for the conditional default probability of a constituent name in terms of the thinning process and the portfolio intensity. This formula is analytically tractable for a wide range of familiar portfolio intensity specifications, including affine and quadratic models. It supports the valuation of securities referenced on the portfolio constituents such as bonds or single-name credit default swaps. Given a specification of the portfolio intensity, we obtain a consistent link among single-name and portfolio derivative prices, which facilitates a joint calibration to single- and multi-name market price data. This brings information contained in the single-name credit swap market to bear on the calibration of a reduced-form model of portfolio loss. The joint calibration is a prerequisite for estimating the sensitivity of the value of a portfolio derivative position to a change in the default risk of a constituent.

We illustrate our results by using random thinning to calibrate a familiar stand-alone model of portfolio loss to market rates of CDX High Yield index, tranche and single-name credit swaps. The thinning processes are specified in terms of a doubly-stochastic matrix of next-to-default probabilities. An element of this matrix represents the probability of a CDX index constituent to be the nth defaulter. We provide a fast and stable quadratic programming algorithm to calibrate this matrix, which guarantees that single-name swap market rates are matched perfectly for a given fit of the portfolio intensity to market CDX rates. The fitted thinning matrix reveals investors' risk-neutral probabilities of ordered defaults, which can be used for the valuation of other derivatives, such as first- or second-to-default swaps referenced on the CDX index.

1.1 Related literature

Stand-alone, intensity-based models of the portfolio loss process can be distinguished by the way in which the intensity is updated as new information arrives.[1] In a doubly-stochastic model the intensity is a function of risk factors that describe the economic state. The factors are exogenous in that they do not have jumps in common with the portfolio loss. Conditional on a realization of the factors, the inter-default times are independent. In a self-exciting model the intensity is a function of risk factors that jump at event times. The intensity responds to an event, reflecting the impact of a default on the surviving firms that is channeled through the complex web of contractual and informational relationships in the economy. Random thinning can be applied to both doubly-stochastic and self-exciting intensity models to generate compatible constituent intensities. It supports a *top-down* alternative to a bottom-up approach in which correlated constituent intensity processes are the modeling primitives.[2]

[1] For alternative specifications, see Arnsdorf & Halperin (2007), Brigo, Pallavicini & Torresetti (2006), Davis & Lo (2001), Ding, Giesecke & Tomecek (2006), Errais, Giesecke & Goldberg (2006), Giesecke & Tomecek (2005), Longstaff & Rajan (2007), Lopatin & Misirpashaev (2007) and Tavella & Krekel (2006).

[2] For bottom-up models, see Duffie & Garleanu (2001), Duffie & Singleton (1998), Giesecke & Goldberg (2004), Jarrow & Yu (2001) and Schönbucher & Schubert (2001), among others.

1.2 Structure of this article

Section 2 formulates a stand-alone reduced-form modeling framework for the portfolio default process that encompasses the specifications proposed in the literature. Section 3 develops random thinning to decompose the default process and shows that a thinning process is a probabilistic model for the next-to-default. The results do not require the existence of an intensity. Therefore, they are applicable to generalized reduced-form models of the default process. Section 4 derives a formula for the constituent default probability, and shows how to evaluate it for many familiar portfolio intensity specifications. A set of examples illustrates the results. Section 5 applies random thinning to calibrate a portfolio intensity model to market rates of CDX index, tranche and single-name credit swaps. Random thinning processes are extracted from the market data. Section 6 concludes. Appendix A discusses a stronger version of random thinning that decomposes the default process as well as its intensity. Appendix B contrasts our single-name default probability formulae with standard alternatives. Appendix C demonstrates that two familiar specifications of a doubly-stochastic model disagree.

2 Default process

Consider a strictly increasing sequence of stopping times (T^n) defined on a complete probability space $(\Omega, \mathcal{F}, P)$ equipped with a right-continuous and complete filtration $\mathbb{F} = (\mathcal{F}_t)_{t\geq 0}$. The T^ns represent ordered default times in a portfolio of credit-sensitive securities such as loans, bonds or credit swaps. They generate the *default process* N given by

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

We assume that N starts at 0 and is non-explosive. It can be decomposed into a sum of a local martingale and a nondecreasing predictable *compensator* A. Giesecke (2007) shows that if the times T^n are totally inaccessible and $\exp(A_T)$ is square integrable for some horizon $T > 0$, then the Fourier transform of the default process N can be expressed in terms of a Laplace transform of A. More precisely,

$$E\left[\exp\left(-\int_t^T r_s ds\right) e^{izY+iv(N_T-N_t)} \,\Big|\, \mathcal{F}_t\right] = \mathcal{E}_t^v\big(\Psi(v), z, T, r, Y\big) \tag{1}$$

for real z, v and $t \leq T$, where r is a stochastic discount rate process, $Y \in \mathcal{F}_T$ is a random payoff at the horizon T, $\Psi(v) = 1 - \exp(iv)$ is the characteristic exponent of the Poisson process and, taking u be a complex number with nonnegative real part,

$$\mathcal{E}_t^v\big(u, z, T, r, Y\big) = E^v\left[\exp\left(-\int_t^T r_s ds\right) e^{izY-u(A_T-A_t)} \,\Big|\, \mathcal{F}_t\right] \tag{2}$$

where the expectation is taken under the equivalent complex measure P^v on $\mathcal{F}_T$ defined by the Radon-Nikodym derivative $\exp(ivN_T + \Psi(v)A_T)$. The Laplace transform (2) is a familiar expression in the credit risk literature. If $A = \int_0^\cdot \lambda_s ds$ is given in terms of an *intensity*[3] λ that represents the conditional default rate, then formula (2) is analogous to the price at t of a security that pays $\exp(izY)$ at T if the issuer survives to T and 0 otherwise, assuming the issuer defaults at intensity $u\lambda$. The calculation of this price is well understood for a wide range of parametric models for λ. Therefore, the counting process transform formula (1) supports a portfolio modeling approach in which the process λ is the primitive. Transform inversion leads to computationally tractable pricing relations for many portfolio derivatives. It also generates explicit expressions for risk measures of debt portfolios. Below, we show how to extend the reach of this approach beyond the portfolio level to include the single-name constituents. Again, the Laplace transform (2) plays a key role: it leads to convenient expressions for constituent default probabilities.

3 Disintegrating the default process

Hedging of single-name exposures associated with a position in a portfolio derivative requires that the multi-name model be calibrated to the single-name market. Risk management demands appropriate capital allocation to sub-portfolios. We propose *random thinning* to facilitate these applications.

3.1 Sub-portfolios and information

Suppose the default process N associated with a portfolio of credit-sensitive securities has compensator A relative to a filtration $\mathbb{F}$. Consider a finite sub-portfolio whose constituents are indexed by a set S and let $\tau(S)$ denote the set of default times τ^k corresponding to firms $k \in S$. The sub-portfolio default process N^S counts defaults in S. It is given by[4]

$$N_t^S = \sum_{k \in S} 1_{\{\tau^k \le t\}}.$$

The sub-portfolio default process need not be adapted to the multi-name model filtration $\mathbb{F}$. While $\mathbb{F}$ must always be fine enough to guarantee that the default process N is adapted, it may not be rich enough to distinguish the identity of a defaulter. In this case, a default event is observable but the identity of a defaulter is not. This simplifying assumption underlies many of the stand-alone reduced-form models of the portfolio loss process in the literature. The assumption is standard because it facilitates the specification of parsimonious intensity models that lead to computationally tractable valuation relations for portfolio derivatives via the default process transform (1).

[3] We do not require that the intensity is a priori predictable. Given a general intensity λ, its predictable projection λ^P with respect to $\mathbb{F}$ generates the same compensator since $\int_0^t (\lambda_s - \lambda_s^P) ds$ is a bounded variation continuous martingale and hence constant. The projection is unique up to indistinguishability. See Dellacherie & Meyer (1982, Chapter VI) and Revuz & Yor (2005, Chapter IV, Theorem 1.2).

[4] If the sub-portfolio contains a single name k only, we use the superscript k instead of $\{k\}$.

To cover the case where the sub-portfolio default process N^S is not adapted, we consider the *optional projection* of N^S onto $\mathbb{F}$.[5] The optional projection is an adapted process that is unique up to indistinguishability. Its value at t is equal to $E[N_t^S \mid \mathcal{F}_t]$ almost surely. Note that if N^S is adapted, then it is equal to its optional projection. Since N^S is nondecreasing, its optional projection is a submartingale. Therefore the optional projection has a Doob-Meyer decomposition into a sum of a nondecreasing predictable process A^S and a martingale.[6] Below we explain the relationship between A^S and the compensator A of the default process of the full portfolio.

3.2 Random thinning

In Proposition 3.1, we express A^S as a Stieltjes integral with respect to the compensator A of the *next-to-default process,* which is denoted Z^S, and defined as an almost sure limit with respect to an appropriate measure on the product space $\Omega \times \mathbb{R}^+$:

$$Z^S = \lim_{\epsilon \to 0} Z^S(\epsilon). \tag{3}$$

The variable $Z_t^S(\epsilon)$ is the conditional probability at time t that the next defaulter is in the sub-portfolio S, given that a default occurs in the ambient portfolio by time $t + \epsilon$:

$$Z_t^S(\epsilon) = \sum_n \frac{P\big[T^n \in \tau(S) \text{ and } T^n \leq t + \epsilon \,|\, \mathcal{F}_t\big]}{P\big[T^n \leq t + \epsilon \,|\, \mathcal{F}_t\big]} 1_{\{T^{n-1} < t \leq T^n\}}. \tag{4}$$

In formula (4), $\epsilon > 0$ and the quotients on the right are taken to be 0 when the denominator vanishes. In particular, $Z^S(\epsilon)$ vanishes whenever all firms in the portfolio are in default. The relevant measure on the product space $\Omega \times \mathbb{R}^+$ is denoted μ_A. As suggested by the notation μ_A is model dependent—it is determined by the compensator A. It is supported on the maximal subset for which there is a positive probability of imminent default.[7] [8]

The measure μ_A is constructed in terms of the random measure ν on $\mathbb{R}^+$ generated by the compensator A.[9] The random measure ν induces a measure μ on $(\Omega \times \mathbb{R}^+, \mathcal{F} \otimes \mathcal{B}(\mathbb{R}^+))$ defined by

$$\mu(X) = E\left[\int_0^\infty X_s \, d\nu_s\right] \tag{5}$$

[5] See Dellacherie & Meyer (1982, Numbers 43–44) for further details.

[6] The sub-portfolio compensator A^S is the dual predictable projection of N^S onto $\mathbb{F}$.

[7] The statement that for all t, the limit (3) exists P-almost surely is inadequate to our needs since it is consistent with infinitesimal operator being ill defined on a set of positive μ_A measure, i.e., a set on which imminent default can occur with positive probability. This is ruled out by the stronger statement proved in Proposition 3.1.

[8] A precise construction of μ_A is given in Dellacherie & Meyer (1982, Chapter VI, Number 86).

[9] Background on random measures is in Dellacherie & Meyer (1982, Chapter VI, Numbers 63–89).

for X non-negative and $\mathcal{F}\otimes\mathcal{B}(\mathbb{R}^+)$-measurable. In many cases of economic importance, μ_A is in the same measure class[10] as the product measure $P \times \lambda m$ where m denotes Lebesgue measure on $\mathbb{R}^+$.[11]

Proposition 3.1. *For any sub-portfolio S, the next-to-default process Z^S is predictable. Furthermore, Z^S is finite and unique almost surely μ_A. The sub-portfolio compensator A^S can be expressed as a Stieltjes integral*

$$A_t^S = \int_0^t Z_s^S \, dA_s. \tag{6}$$

Proof. Since the random measure ν generated by A^S is absolutely continuous with respect to the random measure μ_A, it follows that μ_{A^S} is absolutely continuous with respect to μ_A. The Radon-Nikodym derivative $d\mu_{A^S}/d\mu_A$ is a predictable process and it is shown in Dellacherie & Meyer (1982, Chapter VI, Theorem and Remark 68) that the process A^S is indistinguishable from

$$\int \frac{d\mu_{A^S}}{d\mu_A} \, dA.$$

To complete the proof, we must show that almost surely μ_A, the Radon-Nikodym derivative $d\mu_{A^S}/d\mu_A$ is equal to the infinitesimal operator:

$$\frac{d\mu_{A^S}}{d\mu_A} = \lim_{\epsilon \to 0} Z^S(\epsilon).$$

This is demonstrated with an extension of the well known martingale proof of the Radon-Nikodym theorem in Meyer (1966, Pages 153–154). The extension, given in Airault & Föllmer (1974, Theorem 4.7), expresses the Radon-Nikodym derivative $d\mu_{A^S}/d\mu_A$ as a limit of a martingale defined on a directed set. More precisely, the sums

$$\sum_{l=0}^{k-1} Z_{t_l}^S (t_{l+1} - t_l) 1_{\{t_l < t \le t_{l+1}\}}$$

where $(t_0, t_1, \ldots, t_k)$ is a dyadic partition of $[0, \infty)$ converge μ_A almost surely to $d\mu_{A^S}/d\mu_A$ as the partitions refine. It follows that for μ_A almost all $(\omega, t) \in \Omega \times \mathbb{R}^+$, if t can be expressed as the intersection of dyadic intervals $(t_i, t_i + d_i)$, then

$$\lim_{i \to \infty} Z_{t_i}^S (d_i)(\omega) = Z_t^S(\omega).$$

This, combined with the fact that outside an evanescent set,[12] $Z_t^S(\epsilon)$ can be arbitrarily well approximated by $Z_s^S(d)$ with s and d dyadic, yields the result. □

[10] Two measures are in the same class if they identify the same sets as null.

[11] Here are two examples. If portfolio defaults follow a Poisson process with constant intensity λ, then the measure μ_A is the product measure $P \times \lambda m$. If portfolio defaults are governed by an intensity that is strictly positive and stochastic, then the compensator A is a strictly increasing process and the measure μ_A is equivalent to the product measure $P \times m$ but it is not, itself, a product measure.

[12] A subset of $\Omega \times \mathbb{R}^+$ is evanescent if its projection on Ω has P measure zero. An evanescent set has μ_A measure zero.

Definition 3.2. *A thinning process for the sub-portfolio S is a predictable process Z^S that satisfies formula (6).*

Proposition 3.1 shows that a thinning process exists for every sub-portfolio S and that it can be expressed almost surely μ_A as the probability that the next default is in S, given that default is imminent. In Section 5 we show how to extract these next-to-default probabilities from market rates of index, tranche and single-name swaps. Lemma 3.3 below lists elementary properties of thinning processes that are relevant for the specification of a parametric thinning process. These properties follow immediately from Proposition 3.1 and formulae (3) and (4).

Lemma 3.3. *Thinning processes satisfy the following properties almost surely μ_A:*

(1) Z^S takes values in the unit interval $[0, 1]$.

(2) For any collection S_k of distinct sub-portfolios whose union is the full portfolio, $\sum_k Z^{S_k} = 1$ on $\{(\omega, t) : t \leq \sup_n T^n(\omega)\}$

(3) $Z^S = 0$ on $\{(\omega, t) : t > \sup_n T^n(\omega)\}$

(4) If N^S is adapted then $Z^S = 0$ on $\{(\omega, t) : t > \sup_{k \in S} \tau^k(\omega)\}$.

Property (3) in Lemma 3.3 says that whenever all the names in the portfolio are in default, the thinning process for any sub-portfolio must vanish. Property (4) refines property (3) in case the multi-name filtration $\mathbb{F}$ is rich enough to distinguish the identity of a defaulter. It says that the thinning for S must drop to zero once all the firms in S have defaulted. If the sub-portfolio default process is not adapted, then unless the whole portfolio is in default there is generally a positive probability that a firm in S remains viable and hence the thinning process Z^S is not guaranteed to drop to zero.

Lemma 3.4. *The following statements are equivalent:*

(1) The portfolio default process N admits an intensity λ.

(2) Each sub-portfolio default process N^S admits an intensity given by $Z^S \lambda$.

(3) Each firm default process N^k admits an intensity given by $Z^k \lambda$.

Proof. If λ is an intensity of N, then $A_t = \int_0^t \lambda_s ds$ so Proposition 3.1 implies that

$$A_t^S = \int_0^t Z_s^S \, dA_s = \int_0^t Z_s^S \lambda_s \, ds. \tag{7}$$

It follows that $Z^S \lambda$ is an intensity of A^S so that (1) $\rightarrow$ (2). Statement (3) is a special case of statement (2), and (3) $\rightarrow$ (1) is a consequence of the fact that the intensity of N is given by the sum over constituents k of the single-name intensities, see Proposition 4.4 below, which gives the argument in case $\mathbb{F}$ is fine enough to distinguish the defaulter identities. The general case follows by the same argument, noting that N is equal to the sum over constituents k of the projections of the constituent default processes N^k. □

In statement (1) of Lemma 3.4 the term "intensity" refers to the conditional default rate in the whole portfolio, i.e. the process λ such that $A = \int_0^\cdot \lambda_s ds$. In statements (2) and (3), the term designates the density of the sub-portfolio compensator with respect to Lebesgue measure, see formula (7). If the identity of a defaulter can be distinguished and the sub-portfolio default process is adapted, then this density is literally the conditional rate of default in the sub-portfolio.

A thinning process allows us to recover the compensator A^S of a sub-portfolio from the compensator A of the ambient portfolio. In general, it cannot be used to recover a sub-portfolio default process N^S from the ambient process N. In the special case that we can thin a default process, we say that Z^S replicates N^S. In Appendix A, we derive necessary and sufficient conditions for replication and discuss implications for default modeling.

4 Constituent analysis

Random thinning disintegrates the portfolio default process into the sum of its constituent processes. This section shows how to implement random thinning.

4.1 Constituent default probabilities

We provide an abstract formula for the conditional default probability of a constituent name in terms of the thinning process and the portfolio intensity. In Section 4.3 we develop analytically tractable expressions for this formula.

Proposition 4.1. *Suppose the portfolio default process N admits an intensity λ. For every constituent name k with thinning process Z^k and $0 < t < T$ we have*

$$P\big[t < \tau^k \leq T \,|\, \mathcal{F}_t\big] = \int_t^T E\big[Z^k_s \lambda_s \,|\, \mathcal{F}_t\big]\, ds \tag{8}$$

so the expectation $E[Z^k \lambda]$ defines the probability density of firm k's default time.

Proof. Let $\widehat{N}^k$ be the optional projection of the constituent default process N^k. For $t > 0$ we have $\widehat{N}^k_t = E[N^k_t \,|\, \mathcal{F}_t]$ almost surely. Since $\widehat{N}^k$ has compensator A^k, the process $\widehat{N}^k - A^k$ is a martingale. These two observations imply that

$$P\big[t < \tau^k \leq T \,|\, \mathcal{F}_t\big] = E\big[\widehat{N}^k_T - \widehat{N}^k_t \,|\, \mathcal{F}_t\big] = E\big[A^k_T - A^k_t \,|\, \mathcal{F}_t\big]$$

for $0 < t < T$. Now formula (8) follows from Proposition 3.1, noting that the compensator A has intensity λ, and Fubini's theorem. □

Formula (8) establishes a link between the price of a multi-name derivative providing protection against multiple correlated defaults in the portfolio, and the price of a single-name derivative providing protection against the default of a portfolio constituent. The former is determined by the default process transform (1) in terms of the portfolio intensity. The latter is determined by the default probability formula (8) in terms of the portfolio intensity and the thinning process, assuming the loss at default is independent of the default time. This link facilitates the calibration of a portfolio intensity model to market rates of portfolio *and* single-name derivatives. The joint fit brings the information contained in the single-name credit swap market to bear on the calibration of the portfolio intensity. It is a prerequisite for estimating the sensitivity of the value of a portfolio derivative position to a change in the default risk of a constituent.

The structure of formula (8) differs from that of the classical reduced-form formulae for single-name default probabilities established in Lando (1998), Duffie, Schroder & Skiadas (1996) and Collin-Dufresne, Goldstein & Hugonnier (2004) under different sets of assumptions. The classical formulae are based on an alternative definition of the intensity. Appendix B contrasts this alternative definition with our definition and illuminates the source of the distinction between the default probability formulae.

4.2 Illustrative examples

Before we show how to evaluate the single-name default probability formula (8), we illustrate the specification of thinning processes and the calculation of single-name default probabilities in a simple Poisson model for two alternative filtrations. In both cases, the thinning processes are constant between arrivals. They generate single-name intensities that are updated at events.

Example 4.2 (Defaulter identities can be distinguished)**.** Consider a two credit portfolio whose default process N is a standard Poisson process stopped at T^2 in a filtration $\mathbb{F}$ that is generated by the firm default times τ^1 and τ^2. The compensator is $A_t = t \wedge T^2$ and the intensity is $\lambda_t = 1_{\{t \leq T^2\}}$. It follows that N has distribution given by

$$P\big[N_T = n\big] = \begin{cases} e^{-T} & n = 0 \\ Te^{-T} & n = 1 \\ 1 - e^{-T} - Te^{-T} & n = 2. \end{cases}$$

Let q be a (2×2) doubly-stochastic matrix, i.e. a matrix for which all entries are nonnegative and all rows and all columns sum to 1. The entry q^{kn} of q is the probability that firm k is the nth defaulter. Letting $\Omega^{kn} = \{\tau^k = T^n\} \in \mathcal{F}_{T^n}$, the thinning processes are

$$Z^k_t = \begin{cases} q^{k1} & t \leq T^1 \\ 1_{\Omega^{k2}} & T^1 < t \leq T^2 \\ 0 & T^2 < t. \end{cases}$$

Each constituent intensity $\lambda^k = Z^k\lambda$. Note that Z^k and therefore λ^k drop to zero as soon as firm k defaults, as illustrated in Figure 1. Since the inter-arrival times T^1 and $T^2 - T^1$ are independent and standard exponential, Proposition 4.1 implies that

$$P\big[\tau^k \leq T\big] = \int_0^T E\big[Z^k_s \lambda_s\big]\, ds = 1 - e^{-T} - (1 - q^{k1})Te^{-T}. \tag{9}$$

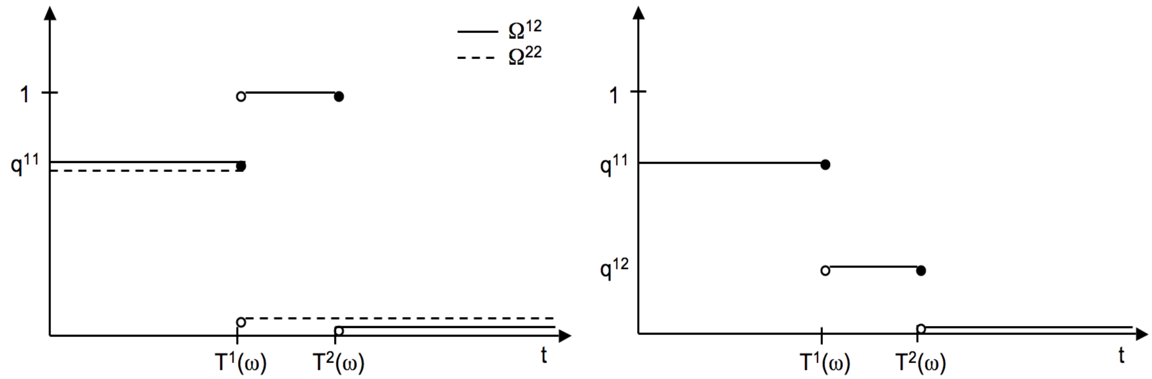

Figure 1: Intensity of firm 1 in Examples 4.2 (left panel) and 4.3 (right panel). The default process is a standard Poisson process stopped at the second default in a filtration that is generated by the firm default times or the default process itself, respectively. The thinning processes are constant between events.

This shows that each τ^k is a mixture of independent exponential variables. If $q^{k1} = 1$, then $\tau^k = T^1$ almost surely so τ^k has the exponential distribution. If $q^{k1} = 0$, then $\tau^k = T^2$ almost surely so τ^k has the gamma distribution. The dependence structure of the default times τ^k is governed by the parameter q^{k1}. The expression for the joint probability of default $P[\tau^1 \leq T, \tau^2 \leq T] = P[N_T = 2] = 1 - e^{-T} - Te^{-T}$ emphasizes the fact that the joint default risk of the two firms is completely determined by the specification of the portfolio intensity λ. The thinning only attributes that risk to the constituents; here this attribution is governed by the parameter q^{k1}. $\square$

In Example 4.2 the filtration is rich enough to distinguish the identities of defaulters. As a result, the thinning processes can reflect the identity of the first defaulter from T^1 onwards. In Example 4.3 below, we thin the stopped Poisson process in the coarser filtration generated by the Poisson process itself. This is the minimal filtration that is compatible with the Poisson model. In this case the thinning processes are given by the projection of the thinning processes in Example 4.2 onto the coarser filtration.

Example 4.3 (Defaulter identities cannot be distinguished)**.** Consider a two credit portfolio whose default process N is a standard Poisson process stopped at T^2 in a filtration $\mathbb{F}$ that is generated by the default process. As in Example 4.2, the intensity is $\lambda_t = 1_{\{t \leq T^2\}}$. The thinning processes are given by

$$Z_t^k = \begin{cases} q^{k1} & t \leq T^1 \\ q^{k2} & T^1 < t \leq T^2 \\ 0 & T^2 < t. \end{cases}$$

Each constituent intensity $\lambda^k = Z^k \lambda$. While Z^k and hence λ^k must drop to zero at T^2, it need not drop to zero at T^1 because the events Ω^{kn} are not observable in the filtration generated by N; see Figure 1. Therefore the single-name intensity λ^k follows a process that is different from the process in Example 4.2. However, the unconditional joint and single-name default probabilities agree with those in Example 4.2 since the filtration generated by N and that generated by the τ^k share the trivial sigma field. □

The informational setting of Example 4.3 is representative of many stand-alone portfolio intensity specifications in the literature in which the filtration is not rich enough to distinguish the defaulter identities. Note that in this case, neither the intensity nor the thinning processes can reflect the identity of a defaulter.

It is instructive to compare the top-down approach taken in the examples above with an alternative bottom-up approach, in which the modeling primitives are the constituent default processes N^k and a filtration $\mathbb{F}$ that is fine enough to distinguish the defaulter identities. Each N^k is specified in terms of its compensator A^k, or more commonly in terms of its intensity λ^k. Together, these components specify the portfolio default process.

Proposition 4.4. *The compensator A to the default process N is given by*

$$A = \sum_k A^k. \tag{10}$$

Proof. A is nondecreasing, predictable and starts at zero. Since $N = \sum_k N^k$, the process $N - A$ is the sum of local martingales $N^k - A^k$ and thus a local martingale itself. Hence A is the compensator in the Doob-Meyer decomposition of N. □

Formula (10) implies that if N admits an intensity λ and the N^k admit intensities λ^k, then $\lambda = \sum_k \lambda^k$ in the sense that they generate the same compensators. The following example is studied, from a different perspective, in Collin-Dufresne et al. (2004), Jarrow & Yu (2001) and Kusuoka (1999). It illuminates the connection between bottom-up and top-down model formulations.

Example 4.5. Consider a two credit portfolio in a filtration $\mathbb{F}$ that is generated by the firm default times τ^1 and τ^2. Suppose the firm intensities satisfy

$$\begin{aligned}
\lambda_t^1 &= \big(p^{11} + p^{12} 1_{\{\tau^2 < t\}}\big) 1_{\{t \le \tau^1\}} \\
\lambda_t^2 &= \big(p^{21} + p^{22} 1_{\{\tau^1 < t\}}\big) 1_{\{t \le \tau^2\}}
\end{aligned}$$

for a matrix $p = (p^{kn})$ of parameters with $p^{k1} \ge 0$ and $p^{k1} + p^{k2} \ge 0$. Thus, the intensity of a firm is constant between events. It jumps at the default of the other firm to incorporate the feedback of the event. The jump sensitivity is specified by the parameter p^{k2}. The intensity of the portfolio default process N is given by Proposition 4.4 as

$$\lambda_t = \lambda_t^1 + \lambda_t^2 = \begin{cases} p^{11} + p^{21} & t \le T^1 \\ (p^{21} + p^{22}) 1_{\Omega^{11}} + (p^{11} + p^{12}) 1_{\Omega^{21}} & T^1 < t \le T^2 \\ 0 & T^2 < t \end{cases}$$

The thinning processes Z^k are the same as in Example 4.2. The portfolio intensity is updated at the first event. If the response in the intensity does not depend on the identity of the first defaulter, i.e. if $p^{11} + p^{12} = p^{21} + p^{22} =: b$, then the default process N is a birth process stopped at T^2 in the filtration $\mathbb{F}$. This simplifying assumption amounts to a specification of N in a filtration that does not distinguish the defaulter identities. As a result, we can immediately calculate the probabilities

$$P\big[N_T = n\big] = \begin{cases} \frac{\Gamma(\frac{a}{b-a}+k)}{\Gamma(\frac{a}{b-a})k!}\big(1 - e^{-(b-a)T}\big)^k e^{-aT} & n = 0,1 \\ 1 - P\big[N_T = 0\big] - P\big[N_T = 1\big] & n = 2 \end{cases}$$

where $a := p^{11} + p^{21} > b$ and Γ is the gamma function. If $a = b$, then N is a Poisson process stopped at T^2 with intensity $\lambda_t = a1_{\{t\leq T^2\}}$. If the parameter matrix p is doubly-stochastic, then $a = 1$ and we obtain the model analyzed in Example 4.2 with $q = p$. In the general case without any restrictions on p, the (2×2) doubly-stochastic matrix q of initial probabilities q^{kn} that firm k is the nth defaulter is determined by

$$q^{k1} = \frac{p^{k1}}{p^{11} + p^{21}}$$

and the property that the columns sum to one. □

In a bottom-up model there is no need to disintegrate the portfolio default process since the constituent default processes are the primitives. However, the researcher must build the dependence structure of all names into each of the single-name intensities. In a top-down model, the name dependence structure is specified on the level of the portfolio intensity. The thinning mechanism guarantees that the dependence is compatibly reflected in the constituent intensities.

It is difficult to incorporate a realistic dependence structure into a bottom-up model while maintaining computational tractability of portfolio derivatives valuation. A simplifying assumption that is widely used is that the constituent default times are conditionally independent given the realization of a risk factor that does not have jumps in common with the N^k. Here, dependence among names is due to the sensitivity of firm intensities λ^k to common risk factors. Event feedback is ignored.

Conditional independence or *doubly-stochastic Poisson* models are discussed in more detail in Appendix C. In particular, we show that bottom-up and top-down doubly-stochastic models are inequivalent. The intuition is foreshadowed by Example 4.5: if we take $p^{k2} = 0$, then the default times are independent and we can think of τ^k as the first jump time of a Poisson process with intensity p^{k1}. This choice generates the simplest bottom-up doubly-stochastic model for (N^1, N^2). The corresponding portfolio default process N has intensity $\lambda_t = p^{11}1_{\{t\leq\tau^1\}} + p^{21}1_{\{t\leq\tau^2\}}$. It follows that N is not a Poisson process but a birth process stopped at T^2.

4.3 Evaluating constituent default probabilities

We show how to evaluate the single-name default probability (8) in terms of the Laplace transform $\mathcal{E}_t^v(u,z,T,r,Y)$ of the compensator $A=\int_0^\cdot \lambda_s ds$ to the portfolio default process N, which is defined in formula (2). The Laplace transform can be calculated explicitly for a wide range of parametric models for λ, including affine and quadratic specifications (Giesecke (2007)). It leads to computationally tractable valuation relations for portfolio credit derivatives through the counting process transform formula (1), and single-name credit derivatives through formula (11) below. Section 5 illustrates the procedure.

Proposition 4.6. *Suppose the portfolio default process N admits an intensity λ such that $\exp(\int_0^T \lambda_s ds)$ is square integrable for some $T>0$. If the thinning processes are given by*

$$Z^k=\sum_{j=1}^n M^{kj} 1_{\{N_-=j-1\}}$$

for $k=1,2,\dots,n$, where $M=(M^{kj})$ is an $(n\times n)$ doubly-stochastic matrix of predictable, unit-valued processes, then for $0<t<T$ we have the default probability formula

$$P\big[t<\tau^k\leq T\,|\,\mathcal{F}_t\big]=\int_t^T\sum_{j=1}^n\Big\{G_t(j-1,s,k,j)-G_t(j-2,s,k,j)\Big\}\,ds \qquad (11)$$

where for $t\leq s\leq T$ and $k,j=1,2,\dots,n$, the function $G_t(m,s,k,j)$ vanishes for integers $m<N_t$ and for all other integers m it is given by

$$G_t(m,s,k,j)=\frac{1}{2\pi i}\int_{-\pi}^{\pi}\frac{e^{-ivm}-e^{iv}}{\Psi(v)}\,e^{ivN_t}\,\partial_z\,\mathcal{E}_t^v(\Psi(v),z,s,0,M_s^{kj}\lambda_s)|_{z=0}\,dv. \qquad (12)$$

Proof. For fixed t and s with $t\leq s$, the P-almost surely nondecreasing function defined by the conditional expectation

$$G_t(x,s,k,j)=E\big[M_s^{kj}\lambda_s 1_{\{N_s\leq x\}}\,|\,\mathcal{F}_t\big],\quad x\in\mathbb{R}$$

corresponds to a measure whose Fourier-Stieltjes transform is given by

$$\mathcal{G}_t(v,s,k,j)=\int_{-\infty}^{\infty}e^{ivx}dG_t(x,s,k,j)=E\big[M_s^{kj}\lambda_s e^{ivN_s}\,|\,\mathcal{F}_t\big],\quad v\in\mathbb{R}.$$

Given the integrability hypothesis on λ, we can apply the transform formula (1), which implies that $E[e^{izY+iv(N_s-N_t)}\,|\,\mathcal{F}_t]=\mathcal{E}_t^v(\Psi(v),z,s,0,Y)$ for an integrable random variable $Y\in\mathcal{F}_s$. Setting $Y=M_s^{kj}\lambda_s$ in this formula and differentiating with respect to z, we get

$$\mathcal{G}_t(v,s,k,j)=-ie^{ivN_t}\partial_z\mathcal{E}_t^v(\Psi(v),z,s,0,M_s^{kj}\lambda_s)|_{z=0}.$$

With this expression, formula (12) is obtained by Fourier inversion. □

5 Calibrating random thinning processes

This section illustrates our theoretical results by thinning a familiar stand-alone portfolio intensity model and calibrating the associated thinning processes to market rates of CDX index, tranche and constituent credit swaps. Throughout this section, the reference probability P is a risk-neutral pricing measure relative to a constant interest rate $r > 0$.

5.1 Portfolio intensity

Suppose that under risk-neutral probabilities, the intensity λ of the portfolio default process N evolves through time according to the self-exciting model

$$d\lambda_t = \kappa(c - \lambda_t)dt + \sigma\sqrt{\lambda_t}dW_t + \delta dL_t \tag{13}$$

where $\kappa \geq 0$, $c > 0$, $\sigma \geq 0$, $\delta \geq 0$ and $\lambda_0 > 0$ are constant parameters such that $2\kappa c \geq \sigma^2$, W is a standard Brownian motion, and $L = \sum_{n\geq 0}^{N} \ell^n$ is the portfolio loss process, where ℓ^n is the loss at the nth event. At time T^n, the loss process jumps and so does λ, with jump size equal to δ times the realized loss ℓ^n. The response of λ reflects the impact of the event on the surviving portfolio constituents. After the event the intensity reverts back to the level c exponentially in mean at rate κ, with diffusive fluctuations that are driven by W. The parameter σ is the exposure of the portfolio constituents to the risk factor W. The model filtration $\mathbb{F}$ is the right-continuous and complete filtration generated by W and L. The defaulter identities cannot be distinguished.

The intensity model (13) can be extended along several dimensions. As discussed in Errais et al. (2006), we can include time-dependent coefficient functions, and additional jump and diffusion terms that describe other risk factors. Here we wish to focus on the random thinning, and therefore choose the relatively parsimonious parametrization (13). In addition, we assume that the loss at default ℓ^n is equal to a constant ℓ for all n. Thus, $L = \ell N$ and it suffices to consider the default process N.

A portfolio derivative is a contingent claim on N. To value a portfolio derivative we require the distribution of N under risk-neutral probabilities. This distribution is given in terms of the characteristic function $E[\exp(iv(N_T - N_t)) \,|\, \mathcal{F}_t]$. Formula (1) expresses this characteristic function in terms of the Laplace transform $\mathcal{E}_t^v(u, z, T, 0, Y)$ of the compensator $A = \int_0^\cdot \lambda_s ds$ under the equivalent complex measure on $\mathcal{F}_T$ defined by the Radon-Nikodym derivative $\exp(ivN_T + \Psi(v)A_T)$. From Giesecke (2007),

$$\mathcal{E}_t^v(u, z, T, 0, \lambda_T) = E^v\left[\exp\left(-u\int_t^T \lambda_s ds\right)e^{iz\lambda_T} \,\middle|\, \mathcal{F}_t\right] = \exp\big(\alpha(t) + \beta(t)\lambda_t\big) \tag{14}$$

for the intensity model (13), where $t \leq T$ and the coefficient functions $\beta(t) = \beta(u, z, v, t, T)$ and $\alpha(t) = \alpha(u, z, v, t, T)$ satisfy the complex-valued ordinary differential equations

$$\partial_t\beta(t) = u + \kappa\beta(t) - \frac{1}{2}\sigma^2\beta(t)^2 - e^{iv}(e^{\delta\beta(t)} - 1) \tag{15}$$

$$\partial_t\alpha(t) = -c\kappa\beta(t) \tag{16}$$

with boundary conditions $\beta(T) = iz$ and $\alpha(T) = 0$. These ODEs are solved numerically. The distribution of N is then obtained by Fourier inversion.

Index and tranche swaps are the most important and actively traded portfolio derivatives. An index swap provides protection against the losses associated with all portfolio defaults. A tranche swap offers insurance against losses exceeding a threshold, up to a maximum amount. As shown in Errais et al. (2006), the fair swap rate can be expressed in terms of prices of call options written on the portfolio default count N. These prices are gotten by integrating the option payoff against the distribution of N. They only depend on the parameters of the portfolio intensity process λ.

5.2 Thinning processes

Suppose that under risk-neutral probabilities, the thinning processes take the form

$$Z^k = \sum_{j=1}^{n} M^{kj} 1_{\{N_- = j-1\}}$$

where $k = 1, 2, \ldots, n$ and n is the number of firms in the portfolio, and $M = (M^{kj})$ is an $(n \times n)$ doubly-stochastic matrix of deterministic, unit-valued functions. Thus, the thinning process Z^k is deterministic between arrivals. The corresponding single-name intensity $\lambda^k = Z^k \lambda$ inherits the attributes of the portfolio intensity λ. In particular, λ^k responds to events, reflecting the impact of a default on firm k. Unless the function M^{kj} has jumps, between events λ^k diffuses according to the Brownian motion W, reflecting the exposure of firm k to a common risk factor.

To value a credit derivative referenced on a portfolio constituent k, we require only the risk-neutral default probability $F_t^k(T) = P[t < \tau^k \leq T \,|\, \mathcal{F}_t]$. Proposition 4.6 expresses this probability in terms of the quantities $\partial_z \mathcal{E}_t^v(u, z, T, 0, M_T^{kj} \lambda_T)|_{z=0}$. Since M^{kj} is deterministic, formula (14) gives

$$\begin{aligned}\partial_z \mathcal{E}_t^v(u, z, T, 0, M_T^{kj} \lambda_T)|_{z=0} &= M_T^{kj} \partial_z \mathcal{E}_t^v(u, z, T, 0, \lambda_T)|_{z=0} \\ &= M_T^{kj} \partial_z \exp(\alpha(u, z, v, t, T) + \beta(u, z, v, t, T)\lambda_t)|_{z=0}\end{aligned}$$

where the coefficient functions α and β satisfy the ODEs (15)–(16). The partial derivative can be expressed in terms of ODEs for $\partial_z \alpha$ and $\partial_z \beta$. Thus, the constituent default probability $F_t^k(T) = \sum_{j=1}^n \int_t^T M^{kj}(s) H_t^j(s) ds$ for some function $H_t^j(s)$ that is determined via formula (11) in terms of four ODEs.

A credit default swap (CDS) referenced on firm k provides protection against the loss due to the default of firm k. The fair swap rate $S_t^k(T)$ for maturity T can be expressed in terms of the default probabilities F_t^k. For constant M^{kj}, it satisfies the equation

$$S_t^k(T) B_t(T) = \sum_{j=1}^{n} M^{kj} \left[g_t^j(T) S_t^k(T) + h_t^j(T) \ell^k \right] \tag{17}$$

where $B_t(T) = \sum_{t<t_m\leq T} \exp(-r(t_m - t))c_m$ with c_m the day count fraction for period m and (t_m) the quarterly coupon dates, $g_t^j(T) = \sum_{t<t_m\leq T} \exp(-r(t_m - t))c_m U_t^j(t_m)$, and $h_t^j(T) = \sum_{t<t_m\leq T} \exp(-r(t_m - t))(U_t^j(t_m) - U_t^j(t_{m-1} \vee t))$ for $U_t^j(T) = \int_t^T H_t^j(s)ds$. The variable ℓ^k is the loss at the default of name k. The model spread depends on the parameters of the portfolio intensity process λ through the functions $H^1, \dots, H^n$ and the parameters $M^{k1}, \dots, M^{kn}$ of the thinning process for name k.

5.3 Calibration results

Our calibration proceeds in two steps. First we calibrate the portfolio intensity parameter vector $\theta = (\lambda_0, \kappa, c, \sigma, \delta)$ to a set of market rates of index and tranche swaps referenced on the CDX High Yield portfolio. The CDX.HY is a standard reference portfolio that consists of $n = 100$ names. Market rates for the five year maturity are observed on May 11, 2007. In accordance with market practice, we suppose that the loss at default ℓ is 60%. The risk-free rate of interest r is 5%. Swap premium payments are made quarterly. The valuation date is $t = 0$. Using gradient based methods, we numerically solve the constrained nonlinear optimization problem

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

$$\text{subject to } 2\kappa c \geq \sigma^2$$

where $\Theta = (0, 5] \times [0, 5] \times (0, 5] \times [0, 2] \times [0, 5]$ and the sum ranges over the quotes. The market mid quote MarketMid(m) is the arithmetic average of the market bid MarketBid(m) and ask MarketAsk(m) for index or tranche m. The model rate is denoted Model(m,θ). The index and tranche swap pricer and the optimization were implemented in Matlab. The computations were performed on a desktop PC with a 1.6 GHz Dual Core Intel processor and 2 GB of RAM. The market and fitted rates are shown in Table 1. The model fits the data, with an average absolute percentage pricing error (AAPE) of 0.87%. Three out of five contracts are priced within the bid-ask spread. The optimal parameters are

$$\hat{\theta} = (0.68, 1.88, 1.41, 0.35, 2.49). \tag{18}$$

These values are insensitive to the initial values: we ran 100 distinct optimizations, each based on random initial values drawn from a uniform distribution over the parameter space Θ. For each set of initial values the optimization converged to values very close or equal to the optimal values in (18), indicating the stability of the calibration procedure.

Taking the fitted portfolio intensity parameters (18) as given, in a second step we calibrate the doubly-stochastic matrix M to the rates of the five year single-name credit swaps referenced on the 100 constituents of the CDX High Yield portfolio, all observed on May 11, 2007. We adjust the single-name market swap rates to account for differences in the terms of the index and constituent swap contracts. For example, some types of credit

Portfolio	Contract	MarketBid	MarketAsk	Fitted
	Index	262.85	263.10	262.97
	0-10%	70.50%	70.75%	71.69%
CDX.HY	10-15%	34.25%	34.50%	33.48%
	15-25%	316.00	319.00	317.86
	25-35%	79.00	81.00	79.90
AAPE				0.87%

Table 1: CDX High Yield index and tranche swap calibration data and results. The five year swap market bid and ask rates were observed on 5/11/2007. The index swap pays all default losses in the underlying portfolio of 100 names that occur before the five year maturity. A tranche swap is specified by two percentage attachment points. It pays the default losses exceeding the lower attachment point. The payoff is capped at the upper attachment point. The 0-10 percent and 10-15 percent tranches are quoted in terms of a percentage upfront rate that is paid at contract inception. All other tranches and the index are quoted in terms of a running spread rate that is measured in basis points (10^{-4}) and paid quarterly. We report the average absolute percentage fitting error (AAPE) relative to mid-market quotes.

events that trigger a payoff in a single-name credit swap do not lead to a payoff in the index contract. For this reason, the market index rate may differ from the intrinsic bottom-up index rate, which is given by a weighted average of the single-name rates (Errais et al. (2006)). In accordance with market practice, we determine a common adjustment that applies to every single-name rate such that the corresponding intrinsic index rate matches the observed market index rate.

The market mid quote MarketMid(k) is the arithmetic average of the adjusted market bid and ask quotes for the swap referenced on firm k. Since we are calibrating to a single maturity only, it suffices to take the entries of M to be non-negative constants. Motivated by the swap pricing formula (17), we formulate the single-name calibration problem as the quadratic programming problem

$$\min_{M} \sum_{k=1}^{100} \Big\{ \sum_{j=1}^{100} M^{kj} \big[g^j(\theta) \cdot \text{MarketMid}(k) + h^j(\theta)\ell^k \big] - B \cdot \text{MarketMid}(k) \Big\}^2$$

$$\text{subject to } M = (M^{kj}) \text{ doubly-stochastic}$$

where $g^j(\theta) = g_0^j(\theta, T)$, $h^j(\theta) = h_0^j(\theta, T)$ and $B = B_0(T)$ are defined in Section 5.2 above, T is equal to five years, and the portfolio intensity parameter vector θ is set to the calibrated value (18). The loss ℓ^k at the default of name k is set to 60%, consistent with market practice and the portfolio intensity calibration. We use the Mosek Matlab toolbox to address the quadratic programming problem.

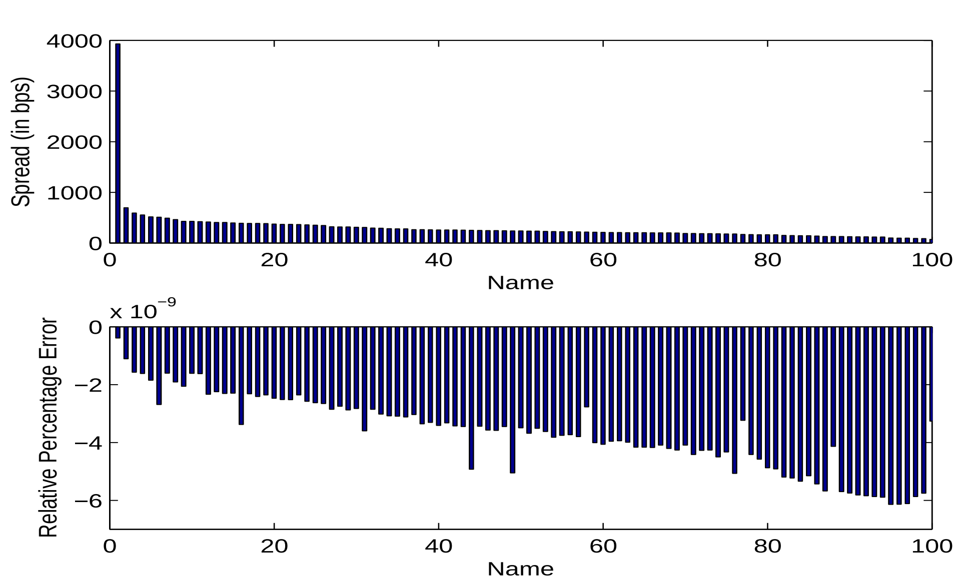

Figure 2: Single-name swap calibration data and results. *Upper panel:* Mid-market single-name credit swap spreads of all 100 CDX High Yield portfolio constituents. The names are ordered according to their spreads. Name 1 is the firm with the lowest creditworthiness (Tembec, Inc.) and name 100 is the firm with the highest credit quality. *Lower panel:* Relative fitting errors, measured in 10^{-9} percent. The fitting errors are essentially zero.

Figure 2 shows the mid-market single-name swap market rates and the fitting errors. All swap rates are matched perfectly. Figure 3 shows the corresponding fitted (100×100) doubly-stochastic matrix M of next-to-default probabilities that represents the thinning processes Z^k for all 100 names. The kth row of M represents the risk-neutral probability of name k to be the jth defaulter, for $j = 1, \ldots, n$. Figure 4 shows these probabilities for the riskiest name in the portfolio (Tembec, Inc.), and the least risky name. As expected, Tembec is assigned the highest probability of defaulting first. The jth column of M represents the risk-neutral probability of name k to be the jth defaulter, for each $k = 1, \ldots, n$. The fitted risk-neutral default probability term structures $\{P[\tau^k \leq T] : 0 < T \leq 5Y\}$ are shown in Figure 5. The fitted term structures and next-to-default probabilities can be used for the valuation of other derivatives, such as first- or second-to-default swaps referenced on the CDX High Yield portfolio.

We can reduce the effective dimension of the doubly-stochastic matrix M. Suppose we focus on the first d columns of M only and distribute mass uniformly over the elements in the remaining columns. This choice is reasonable since higher order next-to-default probabilities are very hard to pin down. The average fitting error across all names is decreasing in d. However, the average fitting error only rises significantly above zero for $d \leq 20$. The errors tend to be higher for high-quality names.

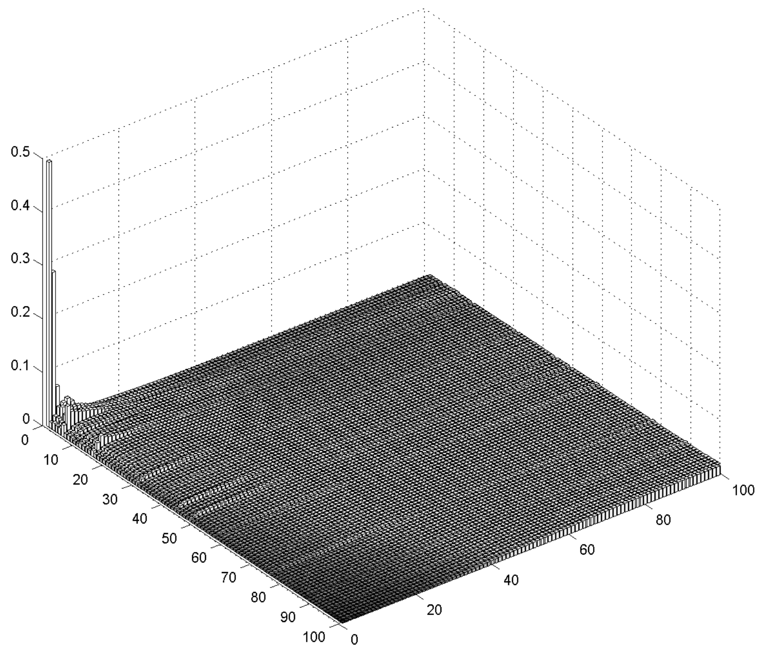

Figure 3: The fitted (100×100) doubly-stochastic matrix M of next-to-default probabilities that represents the thinning processes for all 100 names. The riskiest name 100 (Tembec, Inc.) is assigned the highest probability of defaulting first, roughly 50%.

6 Conclusion

A multi-name credit derivative is a contract tied to an underlying portfolio of corporate bonds or other defaultable assets. Multi-name derivatives play a key economic role: they provide tailored insurance against default losses in the reference portfolio, and therefore allow investors to mitigate their exposure to correlated default risk. A popular approach to analyzing these derivatives is to specify the portfolio loss process without reference to the constituent names, and then to value the contract as a claim written on portfolio loss. This approach leads to a derivative pricing problem in which the underlying "security" follows a point process with given intensity. Several computational tools are available to efficiently address this pricing problem.

In this article, we show how to extend the reach of this methodology to the portfolio constituent securities. Our top-down approach allows the researcher to consistently analyze the portfolio derivative and the securities referenced on its constituents. It facilitates important applications that cannot be addressed with a stand-alone model of the portfolio loss process, for example single-name market calibration or constituent risk hedging. At the center of the top-down approach is random thinning, which allocates the intensity

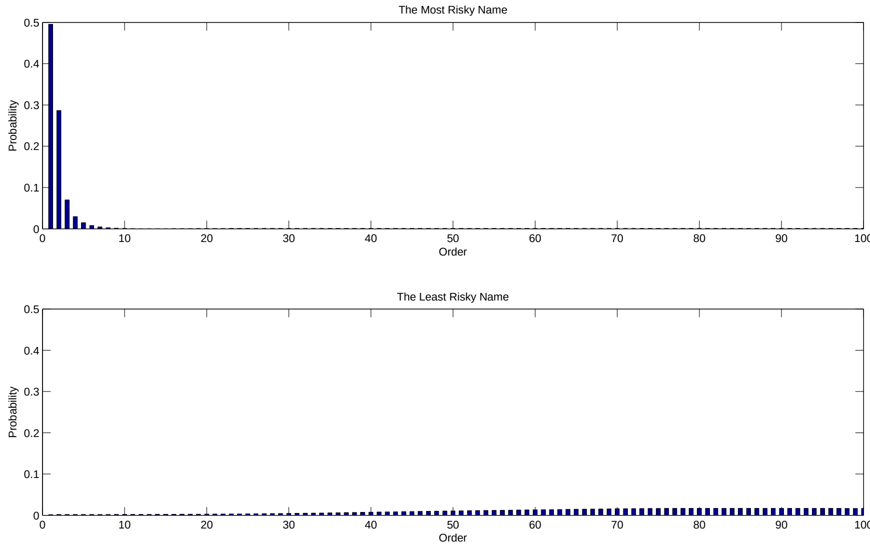

Figure 4: Two rows of the fitted (100×100) doubly-stochastic matrix M of next-to-default probabilities. *Upper panel:* Probability of the riskiest name in the CDX High Yield portfolio, Tembec, to be the jth defaulter, for $j = 1, \ldots, n$. *Lower panel:* Probability of the least risky name in the CDX to be the jth defaulter.

governing the portfolio default process to the constituent single-name default processes. We show that under mild technical conditions, any default process can be thinned, and that the thinning process is a probabilistic model of the next-to-default. We derive a formula for the constituent default probability in terms of the portfolio intensity and the thinning process, and show how to evaluate it for familiar portfolio intensity models. We develop and implement a fast and stable calibration algorithm to extract the thinning processes from market rates of multi- and single-name credit derivatives.

A Strong random thinning

In general, a thinning process Z^S given by Definition 3.2 cannot be used to recover the sub-portfolio default process N^S from the default process N. In this appendix, we consider thinning on the level of default processes. The analysis has implications for structural multi-name modeling and the simulation of intensity-based multi-name models.

Throughout this appendix, we assume that the model filtration $\mathbb{F}$ is fine enough to distinguish the identity of a defaulter. In other words, the constituent default times τ^k are stopping times. This assumption is satisfied in the cases of interest.

Definition A.1. *A strong thinning process for the default process N with respect to a*

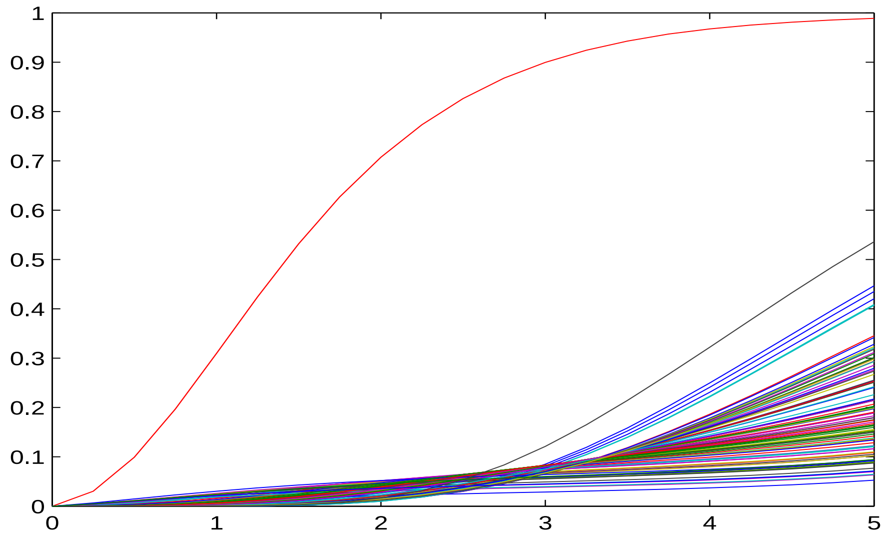

Figure 5: Fitted default probability term structures for all 100 CDX High Yield portfolio constituents. The riskiest name is Tembec, Inc.

stopping time τ is a bounded,[13] *predictable process Z for which the Stieltjes integral*

$$\int_0^t Z_s dN_s \tag{19}$$

is a version of the indicator process of τ.

Proposition A.2. *If Z^k is a strong thinning process for the default process N with respect to the constituent default time τ^k, then the Stieltjes integral*

$$\int_0^t Z_s^k \, dA_s \tag{20}$$

defines a version of the single-name compensator A^k.

Proof. Using formula (19), we get

$$M_t^k = N_t^k - \int_0^t Z_s^k dA_s = \int_0^t Z_s^k dM_s \tag{21}$$

where $M = N - A$. Since M is a local martingale so is M^k. Since A is a nondecreasing and predictable process starting at 0 and Z^k is a non-negative process, the process defined by formula (20) is a nondecreasing and predictable process starting at 0. The uniqueness of the compensator yields the result. □

[13]The assumption that the process Z be bounded can be weakened. For example, it is sufficient that Z be integrable with respect to the local martingale $N - A$. If we make this a priori weaker assumption, Propositions A.2 and A.3 still hold and the boundedness of a strong thinning process is a consequence.

Next, we characterize the situation when a thinning vector is a strong thinning vector. Note first that formula (20) implies that the value of a thinning process is determined only on sets where the compensator A is increasing. Therefore, the values of the thinning vector are well defined only up to a measure μ_A subset of $\Omega \times \mathbb{R}^+$.

Let C denote the set of coordinate vectors, which are unit vectors that have a single coordinate equal to one and all other coordinates equal to zero.

Proposition A.3. *A strong thinning vector* $\mathbf{Z}$ *for the default process* N *takes values in the set* C *of coordinate vectors* μ_A*-almost everywhere.*

Proof. The constituent default process N^k is given by the pathwise Stieltjes integral

$$N_t^k = \int_0^t Z_s^k \, dN_s = \sum_j Z_{\tau^j}^k 1_{\{\tau^j \le t\}} \tag{22}$$

P almost surely. It follows that off a P-measure zero set of paths, the strong thinning vector $\mathbf{Z}$ takes values in C at the event times. Let $W \subset \Omega \times \mathbb{R}^+$ denote the set for which the thinning vector $\mathbf{Z}$ does not take values in C. Then $W \in \mathcal{F}_P$, the predictable sigma algebra on $\Omega \times \mathbb{R}^+$, and it intersects the graphs of the default times on an evanescent set. We show that W has μ_A measure zero.

The counting process N induces a measure μ_N on $(\Omega \times \mathbb{R}^+, \mathcal{F} \otimes \mathcal{B}(\mathbb{R}^+))$ that is analogous to μ_A. Since A is the compensator, or the dual predictable projection of N, Dellacherie & Meyer (1982, Section VI, Theorem 65) implies the two measures μ_A and μ_N agree as linear functionals on non-negative, predictable processes.

Therefore, $0 = \mu_N(1_W) = \mu_A(1_W)$. □

Proposition A.3 says that at any time at which default risk is non-trivial, the strong thinning vector $\mathbf{Z}$ points to a single-name as the next defaulter. Since each coordinate process Z^k of $\mathbf{Z}$ is predictable by hypothesis, it follows that "the name of the next defaulter" is also a predictable process. This is economically unrealistic for two reasons. First, it assumes that investors can anticipate the identity of the next defaulter. Second, it shows that a strong thinning vector exists only if at any given time, a single firm carries all the default risk in the portfolio. This is made precise in the next result.

Proposition A.4. *There exists a strong thinning vector* $\mathbf{Z}$ *for the default process* N *if and only if the single-name compensators* A^k *are mutually singular as random measures on* $\mathbb{R}^+$ *almost surely.*

Proof. This is immediate since the coordinates Z^k of the thinning vector $\mathbf{Z}$ are the Radon-Nikodym derivatives of the compensator A^k with respect to A. □

Remark A.5. *Suppose the default stopping times* τ^k *are predictable. Then the random measures* ν_{A^k} *associated to the single-name compensators* A^k *are Dirac measures concentrated at the times* τ^k*. Since the sequence* (T^n) *of default times is strictly increasing, the* ν_{A^k} *are mutually singular. It follows that a strong thinning vector always exists in a portfolio with predictable default times.*

This observation has implications for multi-name models that are based on predictable default times. The multi-name extensions of the structural models of Merton (1974) and Black & Cox (1976) are prime examples. In these models, a firm defaults when its continuous and adapted firm value process falls below a constant barrier. Therefore, the default times are predictable and investors can anticipate events. Remark A.5 implies that in addition, investors anticipate the identity of the next defaulter. Therefore, these models forecast that short term credit spreads are zero. This prediction is inconsistent with empirical data. The prevalence of positive spreads in the market provides evidence that there are aspects of default that cannot be anticipated by investors.

B Single-name default probabilities

Our default probability formula (8) differs markedly from the standard formulae derived in the literature. In this appendix we compare the two formulations to gain further insights into the probabilistic structure underlying the top-down approach.

Let H be a non-explosive counting process defined on a probability space $(\Omega, \mathcal{F}, P)$ with right-continuous and complete filtration $\mathbb{F}$. Let $\tau = \inf\{t > 0 : H_t > 0\}$ be the first jump time of H. The stopping time τ models the default time of a firm. If H has bounded nonnegative intensity h and the process $V(T)$ defined for fixed $T > 0$ by

$$V_t(T) = E\left[e^{-\int_t^T h_s ds} \,|\, \mathcal{F}_t\right], \quad 0 \leq t \leq T \tag{23}$$

does not jump at τ almost surely, then on $\{\tau > t\}$ the default probability satisfies

$$P\left[t < \tau \leq T \,|\, \mathcal{F}_t\right] = 1 - V_t(T), \quad 0 \leq t \leq T$$

see Duffie et al. (1996). The no-jump hypothesis is violated when h depends on τ as in Example 4.5. It can be relaxed by introducing an absolutely continuous change of measure that puts zero probability on paths for which default occurs before T; see Collin-Dufresne et al. (2004). Since h is the intensity of H, the process defined by $H_t - \int_0^t h_s ds$ is a local martingale. It follows that the process given by

$$1_{\{\tau \leq t\}} - \int_0^{t \wedge \tau} h_s ds \tag{24}$$

is a uniformly integrable martingale. Therefore, the indicator process associated with τ has intensity given by $h_t 1_{\{t \leq \tau\}}$. In the traditional intensity-based credit modeling literature, the process h is the modeling primitive.[14] It is chosen such that the calculation of the expectation (23) is computationally tractable.

[14] In this literature, the process h is often called the "intensity of τ." Duffie & Garleanu (2001) use the term "pre-intensity," which is more appropriate.

In the top-down approach we take a different perspective. The martingale property of the process given by (24) implies the alternative default probability formula

$$P\big[t < \tau \le T \,|\, \mathcal{F}_t\big] = \int_t^T E\big[h_s 1_{\{s \le \tau\}} \,|\, \mathcal{F}_t\big] ds, \quad 0 \le t \le T \tag{25}$$

which is expressed in terms of the intensity of τ. Since it does not involve the auxiliary process $V(T)$, formula (25) does not require the no-jump hypothesis. Compare formula (25) with formula (8). In the top-down approach the portfolio intensity λ and the thinning processes Z^k are the modeling primitives. Assuming the filtration is fine enough to make τ^k a stopping time, thinning establishes the process $Z^k\lambda$ as the intensity of the indicator process associated with the default time τ^k of firm k. Therefore the variable $Z^k_s\lambda_s$ in formula (8) takes the role of the variable $h_s 1_{\{s\le\tau\}}$ in formula (25).

C Bottom-up vs. top-down doubly-stochastic models

Doubly-stochastic Poisson processes are widely used in bottom-up multi-name modeling. In this appendix, we contrast a bottom-up doubly-stochastic model with its top-down counterpart. The analysis has implications for multi-name model design.

C.1 Doubly-stochastic model from the bottom-up

A standard definition of a doubly-stochastic model is cast in terms of the single-name default processes N^k, see Lando (1998) or Duffie & Garleanu (2001). The filtration $\mathbb{F}$ is assumed to be fine enough to distinguish the identity of a defaulter.

Definition C.1 (Bottom-up specification of a doubly-stochastic model)**.** *The firm default process $(N^1, \dots, N^n)$ is doubly-stochastic with respect to a right-continuous, complete filtration $\mathbb{H} \subseteq \mathbb{F}$ if its default times τ^k are conditionally independent with respect to $\mathbb{H}$*[15] *and if there are $\mathbb{H}$-predictable processes h^k so that for each $t \ge 0$, $u > 0$ and k,*

$$P\big[\tau^k > t+u \,|\, \mathcal{F}_t \vee \mathcal{H}_{t+u}\big] = \exp\left(-\int_t^{t+u} h^k_s ds\right) \quad on \quad \{\tau^k > t\}. \tag{26}$$

A bottom-up doubly-stochastic model can be specified by giving the underlying sample space Ω a product structure so that $\omega \in \Omega$ is an ordered pair (ω_1, ω_2). The first factor ω_1 supports a stochastic process X that generates the processes h^k. The process X typically represents firm-specific and portfolio wide default risk factors. The filtration $\mathbb{H} \subseteq \mathbb{F}$ is the right-continuous and completed filtration generated by X. Each ω_1 corresponds to a deterministic $h^k(\omega_1)$ for each firm.

[15] The default times τ^k are called *conditionally independent* with respect to a filtration $\mathbb{H} \subseteq \mathbb{F}$ if, for each $t \ge 0$ and $u > 0$, the events $\{\tau^k > t+u\}$ are independent given the σ-algebra $\mathcal{F}_t \vee \mathcal{H}_{t+u}$.

The product structure points to a theoretical built-in limitation of doubly-stochastic models, which is illustrated by fixing a time t and ω_1 and letting ω_2 vary. Then the values of the firm j default indicator $N^j(\omega) = N^j(\omega_1, \omega_2)$ vary with ω_2. However, the process $h^k(\omega) = h^k(\omega_1)$ does not vary with ω_2 since it is deterministic conditional on $\mathbb{H}$. Consequently, the default of firm j cannot directly affect the intensity of firm k unless N^j is $\mathbb{H}$-adapted.

To extend the reach of the doubly-stochastic framework, one might try to build a doubly-stochastic model whose filtration $\mathbb{H}$ is generated by the default process of a large, important firm. However, this is not possible. The default times cannot be stopping times in a doubly-stochastic filtration $\mathbb{H}$. Hence, in a doubly-stochastic model, the default of one firm cannot immediately affect the intensity of another.

Proposition C.2. *In a doubly-stochastic model supported by a filtration $\mathbb{H} \subseteq \mathbb{F}$, the event $\{\tau^k \le s\}$ is never contained in $\mathcal{H}_s$.*[16]

Proof. Suppose on the contrary that $\{\tau^k \le s\} \in \mathcal{H}_s$ for some k and s. Choose t so that $t < s$ and set $u = s - t$. Then

$$P\big[\tau^k > t + u \,|\, \mathcal{F}_t \vee \mathcal{H}_{t+u}\big] = P\big[\tau^k > s \,|\, \mathcal{F}_t \vee \mathcal{H}_s\big] = 1_{\{\tau^k > s\}}.$$

This contradicts Definition C.1. □

Note that Proposition C.2 does not make use of the assumption of conditional independence. However, the following discussion regarding the joint distribution of the default times τ^k does. Since the τ^ks are totally inaccessible in a doubly-stochastic model, any finite subset of them determines a unique copula C. Recall that $C = C(x_1, x_2, \ldots, x_n)$ is subject to the Fréchet bounds: $C_l \le C \le C_u$, where the lower bound is $C_l = \max\{x_1 + x_2 + \ldots + x_n + 1 - n, 0\}$ and the upper bound is $C_u = \min\{x_1, x_2, \ldots x_n\}$. Totally inaccessible default times τ^k are perfectly positively dependent if their underlying copula is the upper Fréchet bound C_u. When $C = C_u$, the variables are deterministic functions of each other.

Proposition C.3. *A doubly-stochastic model does not admit perfectly positively dependent default times.*

Proof. Suppose, on the contrary, that for $k \neq j$, the copula determined by τ^k and τ^j is the upper bound C_u. Then according to Embrechts, McNeil & Straumann (2001, Theorem 2), there is a random variable Z and nondecreasing functions f^k and f^j so that $\tau^k = f^k \circ Z$

[16] For the stopping time property to fail, it suffices that there is a single s such that $\{\tau^k \le s\}$ is not in $\mathcal{H}_s$. Proposition C.2 establishes a stronger result by showing that there is no time s at which the default event $\{\tau^k \le s\}$ is in $\mathcal{H}_s$.

and $\tau^j = f^j \circ Z$. Fix u and suppose without loss of generality that $\tau^k \leq \tau^j$. Then

$$\begin{aligned} P\big[\tau^j > u\big] &= P\big[\tau^j > u, \tau^k > u\big] \\ &= E\left[P[\tau^j > u, \tau^k > u \,|\, \mathcal{F}_0 \vee \mathcal{H}_u]\right] \\ &= E\left[\exp\left(-\int_0^u (\lambda_s^k + \lambda_s^j) ds\right)\right] \\ &< P\big[\tau^j > u\big], \end{aligned}$$

which is a contradiction. □

C.2 Doubly-stochastic models from the top-down

A top-down specification of an intensity-based doubly-stochastic model is cast in terms of the portfolio default process N. The filtration $\mathbb{F}$ may or may not be fine enough to distinguish the identity of a defaulter.

Definition C.4 (Top-down specification of a doubly-stochastic model)**.** *The default process N is doubly-stochastic with respect to a right-continuous, complete filtration $\mathbb{H} \subseteq \mathbb{F}$ if there is a non-negative $\mathbb{H}$-predictable process h so that for each $t \geq 0$, $u > 0$,*

$$P\big[N_{t+u} - N_t = k \,|\, \mathcal{F}_t \vee \mathcal{H}_{t+u}\big] = \begin{cases} p(k, \int_t^{t+u} h_s ds) & k < n - N_t \\ \sum_{i=n-N_t}^{\infty} p(i, \int_t^{t+u} h_s ds) & k = n - N_t \\ 0 & k > n - N_t \end{cases} \tag{27}$$

where $p(i, \alpha) = e^{-\alpha}\alpha^k / i!$ is the probability that a Poisson distributed variable with parameter $\alpha > 0$ equals $i \in \mathbb{N}$ and n is the number of firms in the portfolio.

In case $n = \infty$, Definition C.4 simplifies to the specification of a doubly-stochastic model given by Duffie (2004). Then the intensity of the default process N in the filtration $\mathbb{F}$ is given by $\lambda = h$, an $\mathbb{H}$-adapted process. In case $n < \infty$, the default process is gotten by stopping an infinite counting process with intensity h at the nth event. The intensity of N in the filtration $\mathbb{F}$ is therefore given by $\lambda = h 1_{\{N \leq n\}}$. It is not $\mathbb{H}$-adapted.

Proposition C.5. *Suppose that the number of firms n is at least 2. Fix a filtration $\mathbb{H} \subseteq \mathbb{F}$ and suppose $\mathbb{F}$ is fine enough to distinguish the identity of a defaulter. A model $(N^1, \ldots, N^n)$ that is bottom-up doubly stochastic (in the sense of Definition C.1) does not generate a model N that is top-down doubly-stochastic (in the sense of Definition C.4).*

Proof. Suppose $(N^1, \ldots, N^n)$ is a doubly-stochastic model in the sense of Definition C.1 with respect to $\mathbb{H}$. There are $\mathbb{H}$-predictable processes h^k so that

$$P\big[\tau^k > u \,|\, \mathcal{F}_0 \vee \mathcal{H}_u\big] = \exp\left(-\int_0^u h_s^k ds\right)$$

for all u. The intensity of the firm default process N^k in the filtration $\mathbb{F}$ is given by $\lambda^k = h^k(1 - N^k_-)$. From Proposition 4.4, the intensity of N in $\mathbb{F}$ is given by $\lambda = \sum_k h^k(1 - N^k_-)$. However, λ cannot be adapted to $\mathbb{H}$ since h^k is positive and $\mathbb{H}$-adapted and N^k_- cannot be $\mathbb{H}$-adapted by Proposition C.2. Since the process h in Definition C.4 must agree with the intensity λ on the set $\{N \leq n\}$, it cannot be $\mathbb{H}$-adapted. Therefore N cannot be doubly-stochastic with respect to $\mathbb{H}$ in the sense of Definition C.4. □

It follows that Definitions C.1 and C.4 generate disjoint classes of models. However, they are similar in the sense that the intensity cannot be updated at event times.

References

Airault, Hélène & Hans Föllmer (1974), 'Relative densities of semimartingales', *Inventiones math.* **27**, 299–327.

Arnsdorf, Matthias & Igor Halperin (2007), BSLP: markovian bivariate spread-loss model for portfolio credit derivatives. Working Paper, Quantitative Research J.P. Morgan.

Black, Fischer & John C. Cox (1976), 'Valuing corporate securities: Some effects of bond indenture provisions', *Journal of Finance* **31**, 351–367.

Brigo, Damiano, Andrea Pallavicini & Roberto Torresetti (2006), Calibration of cdo tranches with the dynamical generalized-poisson loss model. Working Paper, Banca IMI.

Collin-Dufresne, Pierre, Robert Goldstein & Julien Hugonnier (2004), 'A general formula for the valuation of defaultable securities', *Econometrica* **72**, 1377–1407.

Davis, Mark & Violet Lo (2001), Modeling default correlation in bond portfolios, *in* C.Alexander, ed., 'Mastering Risk Volume 2: Applications', Prentice Hall, pp. 141–151.

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

Ding, Xiaowei, Kay Giesecke & Pascal Tomecek (2006), Time-changed birth processes and multi-name credit. Working Paper, Stanford University.

Duffie, Darrell (2004), Credit risk modeling with affine processes. forthcoming, *Journal of Banking and Finance.*

Duffie, Darrell & Kenneth J. Singleton (1998), Simulating correlated defaults. Working Paper, GSB, Stanford University.

Duffie, Darrell, Mark Schroder & Costis Skiadas (1996), 'Recursive valuation of defaultable securities and the timing of resolution of uncertainty', *Annals of Applied Probability* **6**, 1075–1090.

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

Embrechts, Paul, Alexander J. McNeil & Daniel Straumann (2001), Correlation and dependence in risk management, *in* M.Dempster, ed., 'Risk management: value at risk and beyond', Cambridge University Press, Cambridge.

Errais, Eymen, Kay Giesecke & Lisa Goldberg (2006), Pricing credit from the top down with affine point processes. Working Paper, Stanford University.

Giesecke, Kay (2007), The correlation-neutral measure for portfolio credit. Working Paper, Stanford University.

Giesecke, Kay & Lisa Goldberg (2004), 'Sequential defaults and incomplete information', *Journal of Risk* **7**(1), 1–26.

Giesecke, Kay & Pascal Tomecek (2005), Dependent events and changes of time. Working Paper, Stanford University.

Jarrow, Robert A. & Fan Yu (2001), 'Counterparty risk and the pricing of defaultable securities', *Journal of Finance* **56**(5), 555–576.

Kusuoka, Shigeo (1999), 'A remark on default risk models', *Advances in Mathematical Economics* **1**, 69–82.

Lando, David (1998), 'On cox processes and credit risky securities', *Review of Derivatives Research* **2**, 99–120.

Longstaff, Francis & Arvind Rajan (2007), An empirical analysis of collateralized debt obligations. Forthcoming, *Journal of Finance*.

Lopatin, Andrei & Timur Misirpashaev (2007), Two-dimensional Markovian model for dynamics of aggregate credit loss. Working Paper, Numerix.

Merton, Robert C. (1974), 'On the pricing of corporate debt: The risk structure of interest rates', *Journal of Finance* **29**, 449–470.

Meyer, Paul-André (1966), *Probability and Potentials*, Blaisdell, London.

Revuz, Daniel & Marc Yor (2005), *Continuous Martingales and Brownian Motion*, Springer-Verlag, Heidelberg.

Schönbucher, Philipp & Dirk Schubert (2001), Copula-dependent default risk in intensity models. Working Paper, Universität Bonn.

Tavella, Domingo & Martin Krekel (2006), Pricing nth-to-default credit derivatives in the PDE framework. Working Paper, Octanti Associates.