OPERATIONS RESEARCH
Vol. 59, No. 2, March–April 2011, pp. 283–300
ISSN 0030-364X | EISSN 1526-5463 | 11 | 5902 | 0283

DOI 10.1287/opre.1100.0855
© 2011 INFORMS

A Top-Down Approach to Multiname Credit

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

Lisa R. Goldberg
MSCI Barra, Berkeley, California 94704, lisa.goldberg@mscibarra.com

Xiaowei Ding
Morgan Stanley, Purchase, New York 10577, xiaowei.ding@morganstanley.com

A multiname credit derivative is a security that is tied to an underlying portfolio of corporate bonds and has payoffs that depend on the loss due to default in the portfolio. The value of a multiname 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 computationally tractable derivatives valuation problems. However, these models are silent about the portfolio constituent risks. They cannot be used to address applications that are based on the relationship between portfolio and component risks, for example, constituent risk hedging. This paper develops a method that extends these models to the constituents. We use random thinning to decompose the portfolio intensity into a sum of constituent intensities. We show that a thinning process, which allocates the portfolio intensity to constituents, uniquely exists, and 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 develop a semi-analytical transform approach to evaluate it. The formula leads to a calibration scheme for the thinning processes and an estimation scheme for constituent hedge sensitivities. An empirical analysis for September 2008 shows that the constituent hedges generated by our method outperform the hedges prescribed by the Gaussian copula model, which is widely used in practice.

Subject classifications: finance: securities; financial institutions: banks; probability: applications.
Area of review: Financial Engineering.
History: Received June 2008; revisions received August 2009, February 2010; accepted March 2010. Published online in *Articles in Advance* December 30, 2010.

1. Introduction

A multiname or portfolio credit derivative is a financial instrument tied to a portfolio of defaultable assets such as loans or corporate bonds. Its payoff depends on the aggregate loss due to default in the underlying portfolio. An index swap, for example, insures loss in full, while a tranche swap covers loss in excess of a fixed threshold and up to a maximum amount. Portfolio credit derivatives play an important role in the financial market because they allow credit investors such as banks and asset managers to buy or sell default insurance on the underlying portfolio. Yet they are challenging to analyze due to the complex economic phenomena that underpin the joint default risk of the portfolio constituents.

One approach is to formulate a stochastic point process model that describes the time evolution of portfolio loss. The portfolio derivative payoff is a function of the value of the loss process at specified times. Valuing the derivative entails calculating the risk-neutral expectation of the discounted payoff of the loss distribution at each cash flow time. The distribution of the loss process is determined by the arrival intensity, which represents the conditional default rate in the underlying portfolio, and the distribution of the loss magnitudes. The derivative valuation problem is particularly tractable if the intensity dynamics are specified on a stand-alone basis, without reference to the portfolio constituents. This formulation supports the application of various computational techniques, including semi-analytical, transform, simulation, tree, and PDE methods. It also allows the researcher to describe salient empirical phenomena, including the cyclical behavior of corporate default rates and the ripple effects associated with defaults, in terms of a concise set of economically meaningful parameters. The resulting models have been shown to accurately fit market prices of portfolio derivatives.

There is an important catch, however. A stand-alone intensity model of the portfolio loss process is silent about the component risks. Therefore, it cannot be used to address applications that are based on the relationship between the price of the portfolio derivative and the prices of securities referenced on the individual portfolio constituents. One such application is the hedging of the constituent exposures associated with a position in a portfolio derivative. Constituent name hedging requires the sensitivities of the portfolio derivative price with respect to changes

in the prices of the single-name derivatives referenced on the constituents. The sensitivities determine the amount of single-name protection on each portfolio constituent to be bought or sold in order to neutralize portfolio derivative price fluctuations due to small changes in the constituent risks.

This paper shows how to extend the reach of a stand-alone intensity model of the portfolio loss process to the constituents. The method attributes portfolio-level risk to the components and facilitates the applications that are based on the relationship between portfolio and constituent risks. At its center is *random thinning*, which is used to decompose the portfolio-level default intensity into the sum of the constituent intensities. The constituent intensities govern the default point processes of the constituent firms, and, together with the loss magnitudes, determine the prices of single-name credit derivatives referenced on these firms. The thinning mechanism guarantees that these prices are consistent with the prices of multiname derivatives, which are governed by the portfolio-level intensity. The resulting constituent models reflect the dependence among names. They can capture, for example, the cyclical behavior of default rates and the ripple effects emanating from a default event.

Using results of Airault and Föllmer (1974) on relative densities of semimartingales, we show that our method applies to any stand-alone intensity model of the portfolio loss process. A thinning process, which allocates the portfolio intensity to a component, uniquely exists for any intensity specification. We demonstrate that it has a natural probabilistic interpretation. The value of a thinning process represents the conditional probability that a component firm defaults next, given that a default in the portfolio is imminent. This characterization facilitates parametric specification of a thinning process and economic interpretation of its estimated values.

While random thinning can be implemented by simulation, we show how to realize it analytically. We develop a transform-based formula for the default probability of a constituent name in terms of the thinning process and the portfolio intensity. This formula extends to the constituent level, the computational tractability of a stand-alone intensity model of the portfolio loss process. This approach applies to many familiar portfolio intensity specifications, including models with affine dynamics and models based on time changes, and it facilitates the semi-analytical valuation of securities referenced on the portfolio constituents. Given a specification of the portfolio intensity, we obtain a consistent link among single- and multiname derivatives prices. This link is a prerequisite for single-name applications.

To illustrate our method and transform tools, we calibrate the thinning processes from market prices of single-name, index, and tranche swaps referenced on the CDX High Yield index, which is a standard reference portfolio of 100 names. The thinning processes are parametrized through a doubly stochastic matrix. An element of this matrix represents the probability of a constituent to be the nth defaulter. We formulate the matrix calibration problem as a (regularized) quadratic program whose unique solution is found within seconds. Single-name swap market rates are matched perfectly for a given fit of the portfolio intensity to market CDX index and tranche rates.

To demonstrate the utility of our method, we consider the problem of estimating constituent hedge sensitivities for a position in a tranche swap referenced on the CDX High Yield index. We formulate a tractable scheme to estimate these sensitivities, and we assess the resulting hedges for September 2008, a month that witnessed significant volatility due to the default of Lehman Brothers and the collapse of American International Group. According to the daily profit or loss of the hedged tranche position, the hedges generated by our model and method outperform the hedges prescribed by the Gaussian copula model, which is the standard model used in the financial industry.

1.1. Related Literature

Ding et al. (2009) and Halperin and Tomecek (2008) apply the methodology developed in this paper to time-changed linear birth process and time-changed nonlinear death process models of the portfolio loss, respectively. Our method can also potentially be applied to the loss process models of Arnsdorf and Halperin (2008), Brigo et al. (2007), Cont and Minca (2008), Davis and Lo (2001), Errais et al. (2010), Longstaff and Rajan (2008), Lopatin and Misirpashaev (2008), Tavella and Krekel (2006), and others. Zhou (2009) applies the thinning method to analyze highly customized credit derivatives that are referenced on multiple portfolios.

By extending the reach of a stand-alone portfolio intensity model to the constituents, our results facilitate a *top-down* approach to portfolio credit risk. In this approach the portfolio intensity is the modeling primitive, and random thinning is used to generate the constituent intensities. This approach enables the researcher to select and estimate a model for portfolio derivatives market data separately from the models for the constituent name market data. It is an alternative to the bottom-up approach pursued by Duffie and Garleanu (2001), Eckner (2009), Jarrow and Yu (2001), Kou and Peng (2009), Papageorgiou and Sircar (2007), and others. In the bottom-up approach, the modeling primitives are the intensities of the constituent names; their sum is the portfolio intensity. Here, the models for single- and multiname derivatives market data must be specified and estimated jointly. However, this formulation supports, without additional steps, applications that are based on the relationship between portfolio and constituent risks. This includes the estimation of hedge sensitivities; see Chen and Glasserman (2008). The price of this convenience is the high dimensionality of the state space, which can be as high as several hundred in practice. This could reduce the computational tractability of portfolio derivatives valuation. In the top-down formulation of the portfolio

derivative valuation problem, the dimensionality is user-controlled. The dimensionality increases only in the context of single-name applications, which are facilitated by the tools developed in this paper.

1.2. Structure of This Article

Section 2 discusses the probabilistic setting and contrasts alternative formulations of the portfolio derivatives valuation problem. Section 3 develops random thinning to decompose the portfolio-level default intensity. Section 4 derives and illustrates a formula for the constituent default probability. Section 5 addresses the specification of the thinning process in alternative information settings. Section 6 develops transform methods to evaluate the default probability formula for a broad class of models. Section 7 provides pricing relations for single-name, index, and tranche swaps. Section 8 uses market data for these swaps to fit an example model specification. Section 9 estimates and evaluates hedge sensitivities based on the fitted model. Section 10 concludes. There are three technical appendices; two provide proofs, and one analyzes an optimization problem.

2. Preliminaries

Consider a portfolio of m credit-sensitive securities such as loans or bonds, issued by entities that might default on their payment obligations. Issuer $k = 1, 2, \ldots, m$ defaults at a stopping time $\tau^k > 0$, defined on a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with a right-continuous and complete filtration $\mathbb{F} = (\mathcal{F}_t)_{t\geqslant 0}$ that models the information flow.[1] The probability $\mathbb{P}$ is a risk-neutral pricing measure relative to a constant risk-free rate r. We suppose that $\mathbb{P}(\tau^k = \tau^n) = 0$ for distinct k and n. The default indicator process N^k for constituent k is given by $N^k_t = 1_{\{\tau^k \leqslant t\}}$. It is zero before default and jumps to one at default. The financial loss at default is given by a random variable $l^k \in \mathcal{F}_{\tau^k}$.

A portfolio credit derivative is a security with cash flows that depend on the value of the portfolio loss process $L = l^1 N^1 + \cdots + l^m N^m$ at a set of future times. For example, a contract might specify the payment of L_T at a maturity date T. The buyer of this security has protection against losses due to defaults before T. To assess the arbitrage-free value of a portfolio derivative, we require the distribution of the portfolio loss process L. If the loss magnitudes l^k are independent of one another and independent of the default times, an assumption that is standard in the literature and that we also make in this paper, then it suffices to consider the distribution of the *default process*

$$N = N^1 + \cdots + N^m. \tag{1}$$

Because the default process is increasing, it can be decomposed into a sum of a martingale and an increasing predictable process A, called the compensator to N. It is well known that A governs the distribution of N.[2] Thus, to address the portfolio derivatives valuation problem, we need to model the process A and the distribution of the loss magnitudes. Suppose there is an adapted, nonnegative, and right-continuous *intensity* process λ such that almost surely

$$A_t = \int_0^t \lambda_s \, ds < \infty. \tag{2}$$

We can specify A in terms of the constituents, whose indicator processes N^k govern the default process via formula (1). Because an indicator process is increasing, it can be decomposed into a sum of a martingale and an increasing predictable compensator A^k. Consider the specification of A^k in terms of an adapted, nonnegative, and right-continuous intensity process λ^k such that almost surely

$$A^k_t = \int_0^t \lambda^k_s \, ds. \tag{3}$$

The intensity represents the conditional default rate of firm k. The martingale property of $N^k - A^k$ implies that the default probability $\mathbb{P}(t < \tau^k \leqslant s \mid \mathcal{F}_t)$ is equal to $\mathbb{E}(\int_t^s \lambda^k_v \, dv \mid \mathcal{F}_t)$. This shows that, along with the distribution of the loss magnitude l^k, the process λ^k governs the value of a single-name derivative referenced on firm k, which is a security with cash flows that depend on the loss generated by the default of name k. Examples of single-name derivatives include corporate bonds and credit swaps.

Lemma 2.1. *The default process N has intensity λ given by*

$$\lambda = \lambda^1 + \cdots + \lambda^m. \tag{4}$$

Formula (4) reflects the representation (1) of the default process N in terms of the constituent indicators N^k. It implies that the vector process $(\lambda^1, \ldots, \lambda^m)$ governs the distribution of N. It follows that, along with the distributions of the loss magnitudes l^k, this vector process governs the value of a multiname derivative referenced on the portfolio. This observation supports a *bottom-up* approach to single- and multiname derivatives valuation, in which the vector process $(\lambda^1, \ldots, \lambda^m)$ is the modeling primitive.

To account for the dependence among defaults in the portfolio, the components of the vector $(\lambda^1, \ldots, \lambda^m)$ must follow dependent stochastic processes. Between default times, the co-movements of the processes λ^k reflect the exposure of the constituents to common or dependent risk factors. When a name defaults, the corresponding intensity drops to zero, and the sum $\sum_k \lambda^k$ jumps downward. This effect can be counteracted by the response of the intensities of the surviving firms, which reflects contagion in the complex web of contractual relationships in the economy. With impacts differing across firms, the process followed by the sum $\sum_k \lambda^k$ has complicated dynamics. As a result, although the distribution of N is completely determined, its explicit calculation might be challenging

and render the corresponding portfolio derivatives valuation problem intractable. The difficulties can be overcome with models for λ^k that ignore the impact of a default on the surviving names and lead to a doubly stochastic formulation with conditionally independent default times, as shown by Mortensen (2006) and Eckner (2009). While computationally convenient, empirical studies cast doubt on the validity of the conditional independence hypothesis; see Das et al. (2007).

An alternative approach to constructing A is to specify the process λ directly, without reference to the portfolio constituents. This means that the default risk in the portfolio is specified at the aggregate level. This change in perspective reduces the dimensionality of the problem and facilitates the specification of a process λ that is simpler than the process $\sum_k \lambda^k$, yet captures the essential features of portfolio default risk. The reduction of complexity is achieved by choosing dynamics for λ that largely ignore the constituent information carried by the underlying filtration. One could assume, for example, that the intensity λ is driven by a set of risk factors that are common to all firms, and that λ responds to a default by a jump that models the contagious impact of an event. If the jump size does not depend on the defaulting firm, then the identity of a defaulter is irrelevant to the dynamics of λ. This leads to a class of parsimonious and tractable models that fit portfolio derivatives markets well; see Arnsdorf and Halperin (2008), Brigo et al. (2007), Cont and Minca (2008), Davis and Lo (2001), Ding et al. (2009), Errais et al. (2010), Longstaff and Rajan (2008), and Lopatin and Misirpashaev (2008).

However, the missing link between λ and the constituent intensities λ^k creates a dilemma. On one hand, it yields computational advantages for the valuation of portfolio derivatives. On the other hand, the model applications depending on the relationship between the prices of portfolio derivatives and the prices of single-name derivatives referenced on the portfolio constituents are out of reach. These applications include the hedging of the constituent exposures associated with a position in a portfolio derivative. Constituent name hedging requires the sensitivities of the portfolio derivative price with respect to changes in the prices of the single-name derivatives referenced on the constituents. The sensitivities determine the amount of single-name protection on each portfolio constituent to be bought or sold in order to neutralize portfolio derivatives price fluctuations due to small changes in the constituent risks.

To address these applications, we develop a top-down approach that extends the portfolio intensity model λ to the constituents. The problem is to disintegrate λ into the sum (4) of the intensities λ^k that govern the prices of single-name derivatives referenced on the constituents. We accomplish this by random thinning, a mechanism that has traditionally been used for the acceptance/rejection sampling of point processes. We realize random thinning analytically, allowing us to extend to the constituent level the computational tractability offered by a portfolio intensity model.

3. Random Thinning

We take as given a $\mathbb{P}$-integrable intensity λ for the default process N of a fixed portfolio with m constituents. The underlying filtration $\mathbb{F}$, which supports the constituent stopping times, loss magnitudes, and potentially other risk factors relevant for default arrivals, is fixed. Neither $\mathbb{F}$ nor λ need to be specified at this point. Consider

$$Z_t^k(\omega) = \lim_{\epsilon \to 0} \frac{\mathbb{P}(t < \tau^k \leqslant t+\epsilon \mid \mathcal{F}_t)}{\mathbb{E}(N_{t+\epsilon} - N_t \mid \mathcal{F}_t)}(\omega) \tag{5}$$

at points $(\omega, t) \in \Omega \times (0, \infty]$ where the limit exists. The quotient on the right side of Equation (5), which is taken to be zero when the denominator vanishes, represents the conditional probability at time t that the next defaulter is firm k, given that a default occurs in the portfolio after time t and by time $t+\epsilon$. The next result states that a *thinning process* Z^k exists almost surely with respect to a measure μ on the product space $\Omega \times (0, \infty]$. The measure μ is defined on the sigma-field of predictable sets in the product space by $\mu(B \times (t, \infty]) = \mathbb{E}((m - N_t)1_B)$ for $B \in \mathcal{F}_t$; see Airault and Föllmer (1974). This measure serves as our reference measure in the subsequent analysis.[3]

Proposition 3.1. *For any constituent k, the thinning process Z^k uniquely exists μ-almost surely, is predictable, and satisfies μ-almost surely*

$$\lambda^k = Z^k \lambda. \tag{6}$$

The proof of Proposition 3.1, given in Appendix A, does not use the integral representations (2) and (3). It requires only the continuity of the portfolio and constituent compensators, or equivalently, the total inaccessibility of the default times. Thus, we obtain a stronger result: μ-almost surely

$$A^k = \int_0^{\cdot} Z_s^k \, dA_s. \tag{7}$$

Equation (7) shows that Z^k is equal to the Radon-Nikodym derivative of the random measure generated by A^k with respect to the random measure induced by A. Because A^k and A are predictable processes, so is Z^k.

Proposition 3.1 is a counterpart to the aggregation Lemma 2.1 and expresses a constituent intensity in terms of the portfolio intensity and a thinning process. We interpret Proposition 3.1 in the context of acceptance/rejection sampling. Generate portfolio defaults according to the intensity λ. At an event, randomly select the identity of a defaulter from the set of surviving constituents. At the nth event stopping time T^n, name k is drawn with conditional probability $Z_{T^n}^k$. This probability is determined by information accumulated up to but not including time T^n, a requirement that corresponds to the predictability of the process Z^k. This scheme generates a path of the vector indicator process $(N^1, \ldots, N^m)$ whose component N^k has intensity

$\lambda^k = Z^k\lambda$. Rather than using Proposition 3.1 in this simulation context, we explore its analytical implications.

Proposition 3.1 facilitates a top-down approach to the analysis of portfolio derivatives, in which the portfolio intensity λ and the thinning process vector $(Z^1, \ldots, Z^m)$ are the modeling primitives. The modeler has flexibility in selecting an intensity model that represents the available market data for portfolio derivative prices. The choice of λ involves model risk because unless λ has a special structure, the available data are insufficient to uniquely determine λ; see Cont and Minca (2008). When single-name applications are required, then the Z^k can be chosen in a second step. This, too, involves model risk, because each choice of Z^k generates its own constituent intensity λ^k from λ via formula (6). Note that model risk cannot be avoided. It is also present in an alternative bottom-up formulation. Here the modeler selects the vector process $(\lambda^1, \ldots, \lambda^m)$, which determines the process λ via formula (4). In this approach, the model selection for single- and multiname derivatives market data cannot be separated.

The structure of the thinning problem imposes intrinsic constraints on the thinning processes Z^k. These are stated in Lemma 3.2; they follow immediately from Proposition 3.1 and Equation (5). In practice, there might be additional fitting constraints that restrict the set of possible parametrizations of the Z^k; see §8.

Lemma 3.2. *The Z^k satisfy the following properties μ-almost surely*:

(1) $Z^k \geqslant 0$.
(2) $\sum_{k=1}^m Z^k = 1$ *on* $\{(\omega, t)\colon t \leqslant T^m(\omega)\}$.
(3) $Z^k = 0$ *on* $\{(\omega, t)\colon t > \tau^k(\omega)\}$.

Property (2) in Lemma 3.2 guarantees that the thinning processes generate constituent intensities λ^k that satisfy Lemma 2.1. Property (3) ensures that λ^k vanishes when name k is in default.

4. Constituent Default Probabilities

We develop a formula that expresses the default probability of a constituent k in terms of the thinning process Z^k and the portfolio intensity λ, which are the modeling primitives in a top-down approach. Along with the distribution of the loss magnitude l^k, the default probability determines the prices of single-name derivatives referenced on firm k. The thinning mechanism guarantees that these prices are consistent with the prices of portfolio derivatives, which are governed by λ and the distribution of the l^k. The link between constituent and portfolio derivative prices facilitates many applications. In particular, it allows us to calibrate the Z^k from market prices of single-name derivatives and to estimate hedge sensitivities.

Proposition 4.1. *The conditional constituent default probability is given by*

$$\mathbb{P}(t < \tau^k \leqslant T \mid \mathcal{F}_t) = \int_t^T \mathbb{E}(\lambda_s Z_s^k \mid \mathcal{F}_t)\, ds. \tag{8}$$

Formula (8) is a consequence of Proposition 3.1 and the martingale property of the process $N^k - \int_0^{\cdot} \lambda_s^k\, ds$. The structure of formula (8) differs from that of the classical formulae for single-name default probabilities established in Lando (1998), Duffie et al. (1996), and Collin-Dufresne et al. (2004) under different sets of assumptions. These classical formulae are not based on the intensity process λ^k. They are formulated in terms of an "extended intensity" or "pre-intensity" process that agrees with λ^k up to τ^k but does not vanish at the default of k. This is because the pre-intensity process is adapted to a subfiltration of $\mathbb{F}$ in which τ^k is not a stopping time.

We illustrate the specification of the thinning processes and the calculation of single-name default probabilities via formula (8). A simple portfolio intensity model for a two-name portfolio suffices to highlight the relevant issues.

Example 4.2. Consider a two-credit portfolio whose default process N is a standard Poisson process stopped at T^2, the time of the second default in the portfolio. The intensity is $\lambda_t = 1_{\{t \leqslant T^2\}}$ and N has distribution

$$\mathbb{P}(N_T = n) = \begin{cases} e^{-T} & n = 0, \\ Te^{-T} & n = 1, \\ 1 - (1+T)e^{-T} & n = 2. \end{cases}$$

For a nonnegative constant q^{k1}, suppose a thinning process takes the form

$$Z_t^k = \begin{cases} q^{k1} & t \leqslant T^1, \\ 1_{\Omega^{k2}} & T^1 < t \leqslant T^2, \\ 0 & T^2 < t, \end{cases} \tag{9}$$

where $\Omega^{kn} = \{\tau^k = T^n\} \in \mathcal{F}_{T^n}$. Property (2) of Lemma 3.2 demands that (q^{11}, q^{21}) is a stochastic vector. The parameter q^{k1} models the conditional probability at time $t \leqslant T^1$ that firm k is the first defaulter, given a default is imminent. At T^1, the identity of the first defaulter is revealed, and Property (3) of Lemma 3.2 requires that the thinning process of the defaulting name drops to zero. At the same time, the thinning process of the surviving firm jumps to one due to Property (2) of Lemma 3.2. Thus,

$$\lambda_t^k = Z_t^k \lambda_t = \begin{cases} q^{k1} & t \leqslant T^1, \\ 1_{\Omega^{k2}} & T^1 < t \leqslant T^2, \\ 0 & T^2 < t. \end{cases} \tag{10}$$

Because the interarrival times T^1 and $T^2 - T^1$ are independent and standard exponential, Proposition 4.1 implies that

$$\begin{aligned} \mathbb{P}(\tau^k \leqslant T) &= \int_0^T \mathbb{E}(\lambda_s Z_s^k)\, ds \\ &= 1 - e^{-T} - (1 - q^{k1})Te^{-T}. \end{aligned} \tag{11}$$

Formula (11) 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 a gamma distribution. The dependence structure of the τ^k is governed by the parameter q^{k1}. The expression $\mathbb{P}(\tau^1 \leqslant T, \tau^2 \leqslant T) = 1 - e^{-T} - Te^{-T}$ emphasizes the fact that the joint default risk of the two firms is governed by the portfolio intensity λ. Random thinning allocates the joint default risk to constituents. □

The simple model considered in Example 4.2 can be extended in several directions. The portfolio intensity λ in this example is adapted to the filtration generated by N. It ignores any additional information that is present in the underlying filtration $\mathbb{F}$. That information includes the identity of a defaulter. It might also include observations of stochastic risk factors that influence the firms in the portfolio. A more realistic model is obtained by allowing the intensity λ to follow a nontrival stochastic process. Then the thinning process Z^k specified by Equation (9) generates a constituent intensity $\lambda^k = Z^k\lambda$ that moves randomly between arrivals and reflects the fluctuations of the risk factors driving λ. With piecewise constant $Z^k s$, however, the random movements of λ^1 and λ^2 over $[0, T^1]$ are governed by the movements of λ and are therefore perfectly correlated. To facilitate nontrivial correlation, we can assume that the $Z^k s$ follow stochastic processes over $[0, T^1]$. That is, we can replace the constant q^{k1} in Equation (9) by a nonnegative, predictable process (q_t^{k1}) that is driven by a risk factor. Property (2) of Lemma 3.2 requires that (q_t^{11}, q_t^{21}) is a stochastic vector almost surely.

Proposition 4.3 is a step toward calculating the default probability (8) when λ or Z^k follow stochastic processes. It expresses the probability (8) in terms of the quantity

$$\varphi_t(n, z, s, Y) = \mathbb{E}(\exp(-zY)1_{\{N_s - N_t = n\}} \mid \mathcal{F}_t) \tag{12}$$

for certain nonnegative random variables $Y \in \mathcal{F}_s$ and $z \geqslant 0$. We consider the computation of this quantity in §6. At this point we note only that $\varphi_t(n, z, s, Y)$ is also useful for the valuation of portfolio derivatives, because it provides the distribution of N for $z = 0$. Thus, with an appropriate parametrization of the thinning processes, the incremental computational effort for the calculation of single-name default probabilities might be relatively insignificant.

Proposition 4.3. *Suppose the thinning processes take the form*

$$Z_t^k = \sum_{n=1}^{m} M_t^{kn} 1_{\{T^{n-1} < t \leqslant T^n\}} \tag{13}$$

for $k = 1, 2, \ldots, m$ *and predictable processes* $M^{kn} \geqslant 0$ *that satisfy* $M_t^{1n} + \cdots + M_t^{mn} = 1$ *almost surely for* $T^{n-1} < t \leqslant T^n$ *and* $n = 1, \ldots, m$. *Then the default probability*

$$\mathbb{P}(t < \tau^k \leqslant T \mid \mathcal{F}_t) = -\sum_{n=N_t+1}^{m} \int_t^T \partial_z \varphi_t(n - 1 - N_t, z, s, M_s^{kn}\lambda_s)|_{z=0}\, ds. \tag{14}$$

The representation (13) of a thinning process is motivated by Example 4.2.[4] Property (2) of Lemma 3.2 requires that off a set of μ-measure 0, the sum over k of the variables M_t^{kn} must equal 1 for $T^{n-1} < t \leqslant T^n$ and $n = 1, \ldots, m$. Property (3) requires that for $n = 2, \ldots, m$, the M^{kn} take the form $M^{kn} = Q^{kn}(1 - N_-^k)$ for some predictable process Q^{kn}. These requirements, especially the latter, constrain the set of admissible parametrizations for M^{kn}. Take, for example, $Q^{kn} = q^{kn}$ for constants q^{kn}. These constants must satisfy $q^{11} + \cdots + q^{m1} = 1$ and $q^{kn} = 1/(m - n + 1)$ for all k and $n = 2, \ldots, m$. We are left with $m - 1$ free parameters, all in the first column of the matrix $q = (q^{kn})$. Perhaps more inconvenient is the fact that M^{kn} depends on N^k, which might make it difficult to evaluate $\varphi_t(n, z, s, M_s^{kn}\lambda_s)$. We address these issues in the following section.

5. Smoothing the Thinning Process

Sections 3 and 4 assume that the constituent default times τ^k are stopping times. If τ^k is a stopping time, the kth thinning process must be zero for $t > \tau^k$ (Lemma 3.2). Contrast this with a hypothetical situation in which the filtration distinguishes the ordered default times T^n but not the identity of a defaulter. This means the time of a default is observable, but the name of the defaulter is not. In this situation, a thinning process does not vanish at a default, simply because the identity of a defaulter is not observable in the filtration. In this section we generalize the analysis in §§3 and 4 by allowing for the possibility that the identity of a defaulter is not observable. This enables us to design an alternative, more flexible specification scheme for the thinning processes. It also leads to a computational trick that simplifies the calculation of default probabilities in practice.

Suppose the default process N has intensity $\bar{\lambda}$ relative to a right-continuous and complete filtration $\bar{\mathbb{F}} \subseteq \mathbb{F}$ in which the random times τ^k might not be stopping times.[5] The process $\bar{\lambda}$ can be obtained from the portfolio $\mathbb{F}$-intensity λ, which is the modeling primitive, and which we assume is bounded. It is given by the optional projection of λ onto $\bar{\mathbb{F}}$. The optional projection is a unique adapted process that satisfies $\bar{\lambda}_t = \mathbb{E}(\lambda_t \mid \bar{\mathcal{F}}_t)$ almost surely.[6] It is useful to think about the projection in terms of this conditional expectation and to view $\bar{\lambda}_t$ as an estimate of λ_t given the information represented by $\bar{\mathcal{F}}_t$. Note that $\bar{\lambda} = \lambda$ if λ is adapted to $\bar{\mathbb{F}}$, i.e., if the dynamics of λ ignore the identity of a defaulter. This is a feature of all extant portfolio intensity models that are computationally tractable.

To formulate the thinning problem for $\bar{\lambda}$, we need to clarify the notion of a constituent intensity relative to $\bar{\mathbb{F}}$. A constituent intensity in the sense of §§2–4 might not exist because the default indicator processes N^k might not be adapted. To extend the notion of a constituent intensity to $\bar{\mathbb{F}}$, we consider the optional projection $\bar{N}^k$ of N^k onto $\bar{\mathbb{F}}$, which satisfies $\bar{N}_t^k = \mathbb{E}(N_t^k \mid \bar{\mathcal{F}}_t) = \mathbb{P}(\tau^k \leqslant t \mid \bar{\mathcal{F}}_t)$ almost surely.[7]

If the random times τ^k are stopping times, then the optional projection $\bar{N}^k = N^k$ and the setting reduces to that analyzed in §§2–4. The law of iterated expectations shows that because N^k is increasing, $\bar{N}^k$ is a submartingale. Therefore, it can be decomposed into a sum of a martingale and its compensator $\bar{A}^k$. As a consequence of the representation in formula (3), there is a nonnegative adapted process $\bar{\lambda}^k$ such that $\bar{A}^k_t = \int_0^t \bar{\lambda}^k_s\, ds$ almost surely. Unless N^k is adapted, this process is not an intensity in the classical sense.

An argument detailed in Appendix B, which extends the argument used to prove Proposition 3.1, shows that the portfolio intensity $\bar{\lambda}$ can be disintegrated into the sum over k of the processes $\bar{\lambda}^k$. Consider

$$\bar{Z}^k_t(\omega) = \lim_{\epsilon\to 0} \frac{\mathbb{P}(t < \tau^k \leqslant t+\epsilon \mid \bar{\mathcal{F}}_t)}{\mathbb{E}(N_{t+\epsilon} - N_t \mid \bar{\mathcal{F}}_t)}(\omega) \tag{15}$$

at all points $(\omega, t) \in \Omega \times (0, \infty]$ where the limit exists. Generalizing Proposition 3.1, a predictable thinning process $\bar{Z}^k$ exists uniquely and satisfies

$$\bar{\lambda}^k = \bar{Z}^k \bar{\lambda} \tag{16}$$

$\bar{\mu}$-almost surely. The relevant measure $\bar{\mu}$ is defined on the sigma-field of predictable sets in the product space by $\bar{\mu}(B \times (t, \infty]) = \mathbb{E}((m - N_t)1_B)$ for $B \in \bar{\mathcal{F}}_t$. The thinning process $\bar{Z}^k$ is equal to the predictable projection of the $\mathbb{F}$-thinning process Z^k onto $\bar{\mathbb{F}}$. The predictable projection is similar to the optional projection, with the difference that it is always a predictable process.[8] The processes $\bar{Z}^k$ satisfy the following properties $\bar{\mu}$-almost surely:

($\bar{1}$) $\bar{Z}^k \geqslant 0$.
($\bar{2}$) $\sum_{k=1}^m \bar{Z}^k = 1$ on $\{(\omega, t)\colon t \leqslant T^m(\omega)\}$.
($\bar{3}$) $\bar{Z}^k = 0$ on $\{(\omega, t)\colon t > T^m(\omega)\}$.

Properties ($\bar{1}$) and ($\bar{2}$) mirror Properties (1) and (2) of Lemma 3.2, which treats the case where the constituent default times are stopping times. In particular, Property ($\bar{2}$) guarantees that $\bar{\lambda}^1 + \cdots + \bar{\lambda}^m = \bar{\lambda}$. Property ($\bar{3}$) is different from Property (3) of Lemma 3.2. It requires that an $\bar{\mathbb{F}}$-thinning process vanishes whenever all the names in the portfolio are in default. Property (3) of Lemma 3.2 demands that an $\mathbb{F}$-thinning process vanishes with the default of the corresponding firm. In this sense an $\bar{\mathbb{F}}$-thinning process is smoother than an $\mathbb{F}$-thinning process. This is a key property.

Proposition 4.1 can be extended to obtain a formula for the conditional default probability relative to the filtration $\bar{\mathbb{F}}$. Details are in Appendix B. We have

$$\mathbb{P}(t < \tau^k \leqslant T \mid \bar{\mathcal{F}}_t) = \int_t^T \mathbb{E}(\bar{Z}^k_s \bar{\lambda}_s \mid \bar{\mathcal{F}}_t)\, ds. \tag{17}$$

Because the information sets $\mathcal{F}_t$ and $\bar{\mathcal{F}}_t$ might be distinct, the conditional probability (17) might differ from the conditional probability (8). However, the *unconditional* default probabilities implied by the right-hand sides of Equations (17) and (8) agree because $\bar{\mathcal{F}}_0 = \mathcal{F}_0$. Therefore, the probability density function of τ^k satisfies $\mathbb{E}(Z^k_s \lambda_s) = \mathbb{E}(\bar{Z}^k_s \bar{\lambda}_s)$. Thus, we have two ways of calculating $\mathbb{P}(\tau^k \leqslant T)$ for a given portfolio intensity model $(\mathbb{F}, \lambda)$. We can model the $\mathbb{F}$-thinning process Z^k and then exploit formula (8). The alternative is to obtain $\bar{\lambda}$ from λ and $\bar{Z}^k$ from Z^k by projection and then use formula (17). Here is an example.

Example 5.1. Consider the two-credit portfolio of Example 4.2. In the filtration $\mathbb{F}$, the default process N is a standard Poisson process stopped at T^2. Its intensity $\lambda_t = 1_{\{t \leqslant T^2\}}$. Relative to the filtration $\bar{\mathbb{F}}$ generated by N, N has intensity $\bar{\lambda} = \lambda$ since λ is adapted to $\bar{\mathbb{F}}$. For nonnegative constants $\bar{q}^{k1}$ and $\bar{q}^{k2}$, suppose an $\bar{\mathbb{F}}$-thinning process takes the form

$$\bar{Z}^k_t = \begin{cases} \bar{q}^{k1} & t \leqslant T^1, \\ \bar{q}^{k2} & T^1 < t \leqslant T^2, \\ 0 & T^2 < t. \end{cases}$$

In view of Property ($\bar{2}$), we require that $(\bar{q}^{11}, \bar{q}^{21})$ and $(\bar{q}^{12}, \bar{q}^{22})$ be stochastic vectors. The parameter $\bar{q}^{k1}$ must be equal to the parameter q^{k1} in Equation (9). This is because $\bar{Z}^k$ is the predictable projection of Z^k in (9) and q^{k1} is a constant representing the probability that firm k is the first defaulter. The parameter $\bar{q}^{k2}$ represents the probability that firm k is the second defaulter. Note that this probability is not degenerate as in Example 4.2, because the identity of a defaulter is not observable in the filtration $\bar{\mathbb{F}}$ generated by N. Because the $\bar{q}^{kn}$ are constants, the matrix $\bar{q} = (\bar{q}^{kn})$ must even be doubly stochastic, i.e., all rows and all columns must sum to 1. Then

$$\bar{\lambda}^k_t = \bar{Z}^k_t \bar{\lambda}_t = \begin{cases} \bar{q}^{k1} & t \leqslant T^1, \\ \bar{q}^{k2} & T^1 < t \leqslant T^2, \\ 0 & T^2 < t. \end{cases}$$

Formula (17) implies that

$$\begin{aligned} \mathbb{P}(\tau^k \leqslant T) &= \int_0^T \mathbb{E}(\bar{Z}^k_s \bar{\lambda}_s)\, ds \\ &= (\bar{q}^{k1} + \bar{q}^{k2})(1 - e^{-T}) - \bar{q}^{k2} T e^{-T} \\ &= 1 - e^{-T} - (1 - q^{k1}) T e^{-T}, \end{aligned} \tag{18}$$

where the second line follows from the doubly stochastic property of the matrix $\bar{q}$ and the fact that $\bar{q}^{k1} = q^{k1}$. Formula (18) agrees with formula (11). □

In practice, we require only unconditional default probabilities $\mathbb{P}(\tau^k \leqslant T)$. The observations in this section have useful implications for the calculation of these probabilities for portfolio intensity models λ whose $\mathbb{F}$-dynamics ignore the identity of a defaulter. We reiterate that models with this property are standard in the literature because they lead to

tractable portfolio derivatives valuation relations. We argue that they can also lead to tractable single-name derivatives valuation relations. This is because for a suitable auxiliary filtration $\bar{\mathbb{F}}$, the $\bar{\mathbb{F}}$-intensity $\bar{\lambda} = \lambda$, and because we can elect to directly specify the $\bar{\mathbb{F}}$-thinning processes $\bar{Z}^k$ rather than first modeling the Z^k. That is, we model the projections $\bar{Z}^k$, respecting Properties $(\bar{1})$ to $(\bar{3})$, and apply formula (17) to calculate $\mathbb{P}(\tau^k \leqslant T)$. Because we leave the Z^k unspecified, we do not obtain the single-name intensities λ^k. This is unproblematic, however: applications require the default probabilities but not the intensities themselves. For example, the price of a single-name derivative referenced on name k can be expressed in terms of the default probability and the distribution of the loss at default; see §7.1. If we observe market prices of this derivative, we can calibrate $\bar{Z}^k$ for a given model λ. This is illustrated in §8 and forms the basis for the estimation of hedge sensitivities. The calibrated $\bar{Z}^k$ represents the fitted next-to-default probabilities in the sense of (16).

To appreciate the benefits of this approach, consider the specification

$$\bar{Z}^k_t = \sum_{n=1}^{m} \bar{M}^{kn}_t 1_{\{T^{n-1} < t \leqslant T^n\}}, \tag{19}$$

which is motivated by Example 5.1 and parallels the model (13). Here the $\bar{M}^{kn}$ are nonnegative, predictable processes. Property $(\bar{2})$ above requires that these processes satisfy $\bar{M}^{1n}_t + \cdots + \bar{M}^{mn}_t = 1$ for $T^{n-1} < t \leqslant T^n$ and $n = 1, \ldots, m$. Property $(\bar{3})$ is guaranteed by (19). The set of admissible parametrizations of the thinning process is less constrained than in formulation (13). In particular, a thinning process cannot depend on N^k. As shown in Appendix B, we have the default probability formula

$$\mathbb{P}(\tau^k \leqslant T) = -\sum_{n=1}^{m} \int_0^T \partial_z \varphi(n-1, z, s, \bar{M}^{kn}_s \lambda_s)|_{z=0}\, ds, \tag{20}$$

where $\varphi(n, z, s, \bar{M}^{kn}_s \lambda_s)$ is defined by Equation (12) at $t = 0$, dropping the subscript 0 on φ. This quantity is easier to calculate than the corresponding quantity $\varphi(n, z, s, M^{kn}_s \lambda_s)$ required in formula (14). This is because $\bar{M}^{kn}$ cannot depend on N^k, while M^{kn} always does. Take, for example, each $\bar{M}^{kn}$ to be a deterministic function of time. This parametrization is not permissible for the M^{kn}, but it has practical advantages as we demonstrate below. Then $\bar{Z}^k$ is piecewise deterministic and

$$\partial_z \varphi(n, z, s, \bar{M}^{kn}_s \lambda_s)|_{z=0} = \bar{M}^{kn}_s \partial_z \varphi(n, z, s, \lambda_s)|_{z=0}. \tag{21}$$

Thus, we need only to calculate $\varphi(n, z, s, \lambda_s)$ to evaluate the default probability (20). In the following section we provide a transform approach to calculating this quantity.

6. Example Calculations

We illustrate the calculations required to make our results operational.

6.1. Transform Approach

Our calculations exploit a formula for the Laplace transform of a counting process developed by Giesecke and Zhu (2007). Consider a nonexplosive counting process H with intensity ν such that $\exp(\int_0^s \nu_t\, dt)$ is integrable for a horizon $s > 0$. Then, for $u, z \geqslant 0$ we have

$$\mathbb{E}(e^{-uH_s - zY}) = \mathscr{E}^u(\psi(u), z, s, Y), \tag{22}$$

where $Y \in \mathscr{F}_s$ is a nonnegative random variable, $\psi(u) = 1 - \exp(-u)$ is the characteristic exponent of the Poisson process, and

$$\mathscr{E}^u(v, z, s, Y) = \mathbb{E}^u(e^{-v\int_0^s \nu_t dt - zY}), \quad u, v, z \geqslant 0 \tag{23}$$

is the Laplace transform of the random variables $\int_0^s \nu_t\, dt$ and Y, taken under the equivalent probability measure $\mathbb{P}^u$ on $\mathscr{F}_s$ defined by the density

$$\frac{d\mathbb{P}^u}{d\mathbb{P}} = \exp(\psi(u) A_s - uN_s). \tag{24}$$

The transform (22) can be inverted to obtain the probability distribution of N. In Equation (22), set $u = -\log v$ for $v \in (0, 1)$ and expand the left-hand side into a power series. Differentiation with respect to v shows that for $n = 0, 1, \ldots$

$$\mathbb{E}(\exp(-zY) 1_{\{H_s = n\}}) = \frac{1}{n!} \lim_{v \to 0} \partial_v^n \mathscr{E}^{-\log v}(1 - v, z, s, Y). \tag{25}$$

To take advantage of these observations, as in Examples 4.2 and 5.1, we begin with a counting process H with intensity ν relative to the filtration $\mathbb{F}$ and define the default process N as $H \wedge m$. Then the corresponding portfolio intensity satisfies $\lambda = \nu 1_{\{N < m\}}$. The key quantity (12) used in formula (20) is given by

$$\varphi(n, z, s, Y) = \mathbb{E}(\exp(-zY) 1_{\{H_s = n\}}) \tag{26}$$

for the values $n < m$ that are relevant. Thus, applying formula (25), the computation of $\varphi(n, z, s, Y)$ reduces to the calculation of the Laplace transform (23). The calculation of this transform is aided by the observation that it can be interpreted as the price of a security that pays $\exp(-zY) \in \mathscr{F}_s$ at s when the discount rate is given by $v\nu$. The calculation of this price is well understood for a wide range of parametric models for (Y, ν), including affine and quadratic specifications (Duffie et al. 2000 and Leippold and Wu 2002, respectively). The measure change underlying (23) calls for a simple adjustment to the $\mathbb{P}$-dynamics of ν and does not degrade computational tractability.

6.2. Example Specification

We illustrate with an example specification that underlies our numerical experiments in §§8 and 9. We assume that the counting process H driving the default process N has intensity ν with dynamics

$$d\nu_t = \kappa(c - \nu_t)\, dt + \delta\, dJ_t, \tag{27}$$

where $\kappa \geqslant 0$, $c > 0$, $\delta \geqslant 0$, and $\nu_0 > 0$ are parameters, and $J = lH$ where l is a constant. As discussed above, we take $N = H \wedge m$ so that the portfolio intensity $\lambda = \nu$ on $\{N < m\}$. Thus, at a default, the portfolio intensity jumps by δl. This response reflects the impact of an event on the other constituents. After the event the intensity reverts back to the level c, exponentially at rate κ. The portfolio loss process $L = lN$.

The specification (27) can be extended along several dimensions. As discussed in Errais et al. (2010), we can introduce time-dependent coefficient functions and additional jump and diffusion terms that describe other risk factors. Here we choose a relatively parsimonious model in order to focus on the implications of random thinning. The calculations below easily extend to more general intensity specifications.

It remains to specify the thinning processes. We can follow the strategy outlined at the end of §5, because the portfolio $\mathbb{F}$-intensity λ is adapted to the right-continuous and complete filtration $\bar{\mathbb{F}} \subset \mathbb{F}$ generated by the default process N. This filtration does not distinguish the defaulter identities. To formulate a model for the $\bar{\mathbb{F}}$-thinning processes, we suppose that $\bar{Z}^k$ takes the form (19) with each $\bar{M}^{kn}$ a deterministic function of time. The corresponding constituent default probability is given by formulae (20) and (21). We require only the function $\varphi(n, z, s, \lambda_s)$ for $n < m$.

6.3. Transform Calculation

We calculate the Laplace transform (23) induced by the model (27). This Laplace transform determines the key quantity (26) via formula (25). We focus on the case $Y = \nu_s$, which leads to an expression for $\varphi(n, z, s, \lambda_s)$. Noting that $\lambda_0 = \nu_0$, from the results in Giesecke and Zhu (2007) we get

$$\begin{aligned}\mathscr{E}^u(v, z, s, \nu_s) &= \mathbb{E}^u(e^{-v\int_0^s \nu_t\,dt - z\nu_s}) \\ &= \exp(\alpha(s) + \beta(s)\lambda_0),\end{aligned} \tag{28}$$

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

$$\partial_s \beta(s) = -v - \kappa\beta(s) + e^{-u}(\exp(\delta l \beta(s)) - 1), \tag{29}$$

$$\partial_s \alpha(s) = c\kappa\beta(s) \tag{30}$$

with boundary conditions $\beta(0) = -z$ and $\alpha(0) = 0$. Then, from formulae (26) and (25), and noting that λ agrees with ν on the set $\{N < m\}$, we obtain

$$\varphi(n, z, s, \lambda_s) = \frac{1}{n!}\lim_{v\to 0}\partial_v^n \exp(a(z, v, s) + b(z, v, s)\lambda_0) \tag{31}$$

for $n < m$, where the coefficient functions

$$\begin{aligned} b(s) &= b(z, v, s) = \beta(-\log v, z, 1 - v, s), \\ a(s) &= a(z, v, s) = \alpha(-\log v, z, 1 - v, s)\end{aligned}$$

satisfy the ordinary differential equations

$$\partial_s b(s) = -\kappa b(s) - 1 + v\exp(\delta l b(s)), \tag{32}$$

$$\partial_s a(s) = c\kappa b(s) \tag{33}$$

with boundary conditions $b(0) = -z$ and $a(0) = 0$. Using Faà di Bruno's formula, the right-hand side of Equation (31) can be expressed in terms of a recursive system of ODEs that govern the partial derivatives $\partial_v^n a(z, v, s)|_{v=0}$ and $\partial_v^n b(z, v, s)|_{v=0}$. These ODEs are derived from Equations (32) and (33) and do not generally admit closed-form solutions. An exception is the case $n = 0$. Differentiation leads to a set of ODEs for $\partial_z \varphi(n, z, s, \lambda_s)|_{z=0}$.

6.4. Fourier Methods

We provide an alternative approach to the calculation of $\varphi(n, z, s, \lambda_s)$ and $\partial_z \varphi(n, z, s, \lambda_s)$ that might have computational advantages over formula (31), especially when n is large. We extend the Laplace transform (22) to the complex plane by replacing u with $-iu$ for real u. The Fourier transform thus obtained satisfies

$$\mathbb{E}(e^{iuH_s - z\nu_s}) = \exp(A(s) + B(s)\nu_0), \tag{34}$$

where the coefficient functions

$$\begin{aligned} B(s) &= B(u, z, s) = \beta(-iu, z, \psi(-iu), s), \\ A(s) &= A(u, z, s) = \alpha(-iu, z, \psi(-iu), s)\end{aligned}$$

satisfy the complex-valued ordinary differential equations

$$\partial_s B(s) = -\kappa B(s) - 1 + \exp(iu + \delta l B(s)), \tag{35}$$

$$\partial_s A(s) = \kappa c B(s) \tag{36}$$

with boundary conditions $B(0) = -z$ and $A(0) = 0$. These ODEs are obtained from Equations (29) and (30), which characterize the transform (22) for the model (27).

Next consider the function

$$G(x, z, s) = \mathbb{E}(\exp(-z\nu_s)1_{\{H_s \leqslant x\}}), \tag{37}$$

which is nondecreasing in $x \in \mathbb{R}$ and whose Fourier-Stieltjes transform is given by

$$\mathscr{G}(u, z, s) = \int_{-\infty}^{\infty} e^{iux}\,dG(x, z, s) = \mathbb{E}(e^{iuH_s - z\nu_s}). \tag{38}$$

The function $G(x, z, s)$ can be recovered by Fourier inversion from $\mathscr{G}(u, z, s)$, which is characterized through (34) by the ODEs (35) and (36). The inversion can be implemented efficiently through the fast Fourier transform. Because $G(x, z, s)$ is piecewise constant, we need only recover $G(x, z, s)$ at integer values of x. The jump of $G(n, z, s)$ at an integer value $n < m$ is the value $\varphi(n, z, s, \lambda_s)$. These

values also lead to the first n bins of the distribution of the default process:

$$\mathbb{P}(N_s = n) = \varphi(n, 0, s, 0), \quad n = 0, 1, \ldots, m. \tag{39}$$

The value $\varphi(m, 0, s, 0)$ at the last bin m is given by the sum of the jumps of $G(n, z, s)$ at integer values $n \geqslant m$. Similarly, the function $\mathbb{E}(\nu_s 1_{\{H_s \leqslant x\}}) = -\partial_z G(x, z, s)|_{z=0}$, whose jump at an integer value $n < m$ represents the quantity $-\partial_z \varphi(n, z, s, \lambda_s)|_{z=0}$, can be recovered from its transform $\mathbb{E}(\nu_s e^{iuH_s}) = -\partial_z \mathcal{G}(u, z, s)|_{z=0}$ by Fourier inversion. Note that calculating this last transform reduces to solving a system of ODEs. We have

$$\partial_z \mathcal{G}(u, z, s)|_{z=0} = \mathcal{G}(u, 0, s)(A_z(u, s) + B_z(u, s)\lambda_0),$$

where the functions $B_z(u, s) = \partial_z B(u, z, s)|_{z=0}$ and $A_z(u, s) = \partial_z A(u, z, s)|_{z=0}$ satisfy

$$\partial_s B_z(u, s) = -\kappa B_z(u, s) + \exp(iu + \delta l B(u, 0, s))\delta l B_z(u, s),$$

$$\partial_s A_z(u, s) = \kappa c B_z(u, s)$$

with boundary conditions $B_z(u, 0) = -1$ and $A_z(u, 0) = 0$.

7. Single- and Multiname Derivatives Valuation

We use the results of the previous section to value single- and multiname derivatives referenced on the underlying portfolio and its constituents. The valuation relations are required for the numerical experiments in the subsequent sections.

7.1. Single-Name Credit Swap

A credit swap is a bilateral financial contract that transfers the credit risk of a reference firm from one investor, the protection buyer, to another investor, the protection seller. The protection seller covers the loss due to the default of the reference firm k. For a swap with notional equal to 1 and premium payment dates (t_p), the protection seller's payments have a value at time 0 given by

$$D^k = l \sum_p e^{-rt_p} \mathbb{P}(t_{p-1} < \tau^k \leqslant t_p), \tag{40}$$

assuming that the constant loss at default $l \in [0, 1]$ is paid at the next scheduled premium date following default. The protection buyer makes a stream of premium payments at the t_p until τ^k or the maturity date of the contract, whichever is earlier. The cash flow at t_p is a fraction S^k of the swap notional 1. Let c_p be the day count fraction for the period p, roughly 1/4 for quarterly payments. Neglecting premium accruals, the value at time 0 of the payments is given by

$$P^k(S^k) = S^k \sum_p e^{-rt_p} c_p \mathbb{P}(\tau^k > t_p). \tag{41}$$

The annualized swap rate is the solution S^k to the equation $D^k = P^k(S^k)$. To evaluate that rate, we require only $\mathbb{P}(\tau^k \leqslant s)$ for various values of s. For our example specification in §6.2, this probability is given by formulae (20) and (21). These formulae are based on the function $\varphi(n, z, s, \lambda_s)$ for $n < m$, which is calculated in §§6.3 and 6.4.

7.2. Index and Tranche Swaps

Index and tranche swaps are based on a portfolio of single-name credit swaps. They are bilateral financial contracts that transfer specific aspects of the aggregate credit risk of the underlying portfolio from one investor, the protection buyer, to another, the protection seller. The protection seller pays a specified function of the portfolio loss L. The protection buyer pays a fee to the seller to obtain coverage.

The swaps are based on a portfolio whose m constituent single-name swaps have notional 1, maturity date T, and premium payment dates (t_p). The loss at default $l \in [0, 1]$. The swap is specified by a lower attachment point $\underline{K} \in [0, 1]$ and an upper attachment point $\bar{K} \in (\underline{K}, 1]$. An index swap has attachment points $\underline{K} = 0$ and $\bar{K} = 1$. The swap notional $K = m(\bar{K} - \underline{K})$. The protection seller covers portfolio losses as they occur, given that the cumulative losses are larger than $\underline{K}$ but do not exceed $\bar{K}$. The cumulative payments at time t, denoted U_t, are given by the "call spread"

$$U_t = (L_t - \underline{K}m)^+ - (L_t - \bar{K}m)^+. \tag{42}$$

The value at time 0 of these payments is given by

$$\begin{aligned} D &= \mathbb{E}\left(\int_0^T e^{-rs} dU_s \right) \\ &= e^{-rT}\mathbb{E}(U_T) + r \int_0^T e^{-rs}\mathbb{E}(U_s)\, ds. \end{aligned} \tag{43}$$

The protection buyer receives the loss payments and in return makes premium payments to the protection seller. A premium payment has two parts. The first part is an upfront payment, which is expressed as a fraction F of the tranche notional K. For an index swap, $F = 0$. The second part consists of payments that are proportional to the premium notional I_t, which is given by $m - N_t$ for an index swap and by $K - U_t$ for a tranche swap with $\bar{K} < 1$. Then, with S denoting the running premium rate, the value at time 0 of the premium payments is given by

$$P(F, S) = FK + S \sum_p e^{-rt_p} c_p \mathbb{E}(I_{t_p}). \tag{44}$$

For a fixed upfront rate F, the running spread is the solution S to the equation $D = P(F, S)$. Solving that equation for a fixed S gives a value for the upfront rate F. Formulae (43) and (44) indicate that the swap rates depend only on European call options $\mathbb{E}((L_s - c)^+)$ with various strikes c and maturities s and expectations $\mathbb{E}(N_s)$. For our example specification in §6.2, we need only to calculate the values $\mathbb{E}((N_s - c)^+)$ for various c and s. These values are obtained by integrating the option payoff function against the default count distribution: $\mathbb{E}((N_s - c)^+) = \sum_{c < n \leqslant m} (n - c)\varphi(n, 0, s, 0)$.

8. Market Calibration

To demonstrate the top-down approach numerically, we use market prices of single-name, index, and tranche swaps to calibrate the parameters of the portfolio intensity and thinning processes specified in §6.2, and then estimate hedge sensitivities. We employ a two-step calibration procedure. First we fit the parameter $\theta = (\lambda_0, \kappa, c, \delta)$ of the portfolio intensity model (27) to market rates of index and tranche swaps written on the reference portfolio. Fixing the calibrated θ, we then fit the parameter matrix $\bar{M} = (\bar{M}^{kn})$ of the thinning processes (19) to market rates of single-name credit swaps referenced on the portfolio constituents. This means we pick $\bar{M}$ so that the model implied constituent default probabilities match the probabilities implied by the credit swap market rates. Although it involves a high-dimensional parameter space, we formulate this second optimization problem in such a way that it can be efficiently addressed.

8.1. Portfolio Intensity

Our multiname data, obtained from UBS, consist of market rates of $T = 5$ year index and tranche swaps referenced on the CDX High Yield portfolio. The CDX.HY is a standard reference portfolio that consists of $m = 100$ names of relatively low credit quality. The tranches have attachment points $(0, 10\%)$, $(10\%, 15\%)$, $(15\%, 25\%)$, and $(25\%, 35\%)$.

We recalibrate θ from the CDX.HY Series 10 index and tranche rates on each of the 21 trading days of September 2008. That month witnessed significant volatility due to the default of Tembec Industries on the 4th, the demise of Fannie Mae and Freddy Mac announced on the 7th, the default of Lehman Brothers on the 15th, the collapse of American International Group on the 16th, the problems appearing at Morgan Stanley on the 18th, and the bankruptcy of Washington Mutual on the 26th.

Using a gradient-based method, we numerically solve the problem

$$\min_{\theta\in\Theta} \sum_j \frac{(\text{Mid}(j) - \text{Model}(j,\theta))^2}{\text{Mid}(j)}, \tag{45}$$

where $\Theta = (0, 5] \times [0, 5] \times (0, 5] \times [0, 5]$ and the sum ranges over the market rates on a given day. The market mid-quote $\text{Mid}(j)$ is the arithmetic average of the market bid and ask quote for index or tranche j. The model rate $\text{Model}(j, \theta)$ for index or tranche j is given by the formulae developed in §7.2. Swap premium payments are made quarterly. The risk-free rate of interest r is set to 5%. The loss at default l is fixed at 60%, consistent with standard market practice[9] and most other stand-alone intensity based models of the portfolio loss process in the literature. The optimization is initialized at a set of parameter values drawn from a uniform distribution over the parameter space Θ and is repeated for each of 100 independent draws. The optimal parameter θ^* is the solution to (45) with the minimum objective function value among all 100 runs.

The valuation formulae and the optimization are implemented in Matlab. The computations are based on the method described in §6.4. The ODEs (35)–(36) determining the transform (34) are solved numerically using the Runge-Kutta method. The transform (34) is inverted numerically using the fast Fourier transform. The computations are performed on a PC with a 2.66-GHz Intel processor and 4 GB of RAM.

For each of the 21 calibration dates, we calculate the average absolute percentage pricing error (AAPE), given by $(1/5)\sum_{m=1}^{5} |\,\text{Model}(m, \theta^*) - \text{Mid}(m)\,|\,/\text{Mid}(m)$. The left panel of Figure 1 shows the time series of these errors for September 2008. The model fits the market on each date, with an acceptable average AAPE over all calibration dates of 4.4%, and a variance of 0.086. The calibration errors are relatively low during the most volatile period, which started on September 15 with the default of Lehman Brothers. This indicates that the model captures the default correlation implied by the market prices, especially in periods of extreme stress. The calibrated parameter values are shown in the right panel of Figure 1. The time series behavior of the calibrated parameters clearly reflects the events in mid-September.

8.2. Thinning Processes

The single-name market data, obtained from Bloomberg, consist of market rates of $T = 5$ year credit swaps referenced on each of the $m = 100$ constituents of the CDX High Yield portfolio, for each of the 21 trading days of September 2008.[10] We fit a matrix $\bar{M}$ on each of these days, given the corresponding fitted portfolio intensity parameter θ. It suffices to take the entries of $\bar{M}$ to be constants $\bar{q}^{kn} \geqslant 0$. To fit to contracts with different maturities, we can take the entries of $\bar{M}$ to be constant between the available maturities, similar to the parametrization in Halperin and Tomecek (2008). Equations (40) and (41) imply that the model swap rate S^k satisfies

$$S^k W = \sum_{n=1}^{m} \bar{q}^{kn}(V_\theta^n S^k + X_\theta^n l), \tag{46}$$

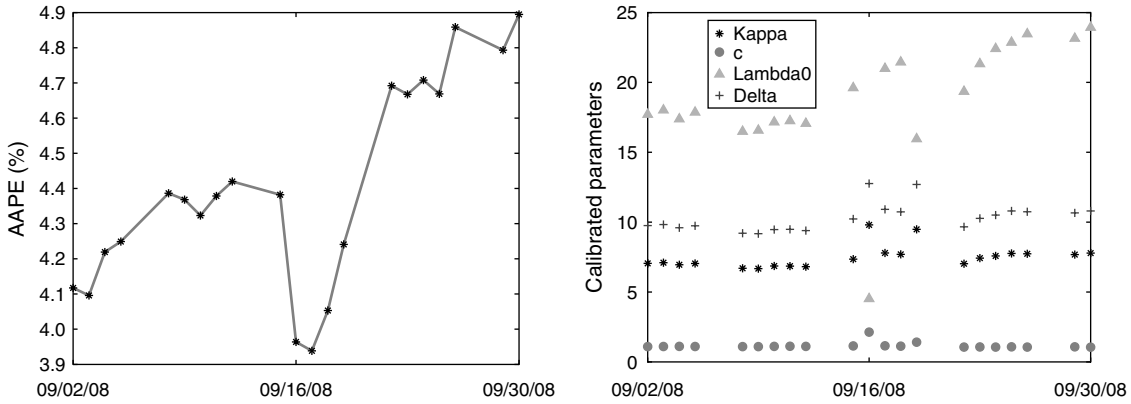
where $W = \sum_p e^{-rt_p} c_p$, and

$$V_\theta^n = -\sum_p e^{-rt_p} c_p \int_0^{t_p} \partial_z \varphi(n-1, z, s, \lambda_s)|_{z=0}\, ds, \tag{47}$$

$$X_\theta^n = -\sum_p e^{-rt_p} \int_{t_{p-1}}^{t_p} \partial_z \varphi(n-1, z, s, \lambda_s)|_{z=0}\, ds. \tag{48}$$

The notation in (47) and (48) indicates the dependence on the parameter θ of the portfolio intensity. The model swap rate will therefore depend on θ and the parameters $\bar{q}^{k1}, \ldots, \bar{q}^{kn}$ of the thinning process $\bar{Z}^k$. Motivated

Figure 1. Calibration results for the portfolio intensity model (27) for the 21 trading days of September 2008, for the CDX High Yield portfolio.



Note. Left panel: Average absolute percentage error (AAPE) for each calibration date. *Right panel:* Value of the calibrated parameter $\theta = (\lambda_0, \kappa, c, \delta)$ for each calibration date.

by relation (46), we formulate the single-name calibration problem as a regularized quadratic programming problem:

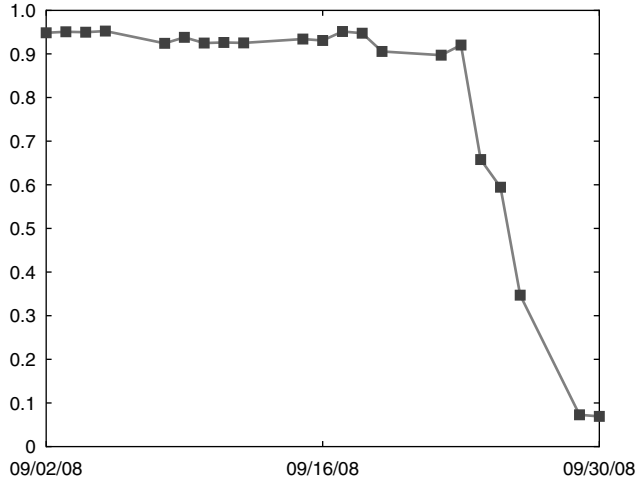
$$\min_{(\bar{q}^{kn})\in\mathcal{M}} \sum_{k=1}^{m} \left(\sum_{n=1}^{m} \bar{q}^{kn} (V^n_{\theta^*} \mathrm{Mid}(k) + X^n_{\theta^*} l) - W\mathrm{Mid}(k) \right)^2 + \sum_{k=1}^{m}\sum_{n=1}^{m} (\bar{q}^{kn})^2, \tag{49}$$

where $\mathcal{M}$ is the set of $(m \times m)$ doubly stochastic matrices, and θ^* is the calibrated value of θ. As explained in Appendix C, the regularization term $\sum_{k=1}^{m}\sum_{n=1}^{m}(\bar{q}^{kn})^2$ guarantees the uniqueness of the solution to the problem. The regularization gives preference to solutions with smaller norms. It affects all entries of $\bar{M}$ uniformly and tends not to distort the structure of $\bar{M}$ imposed by the data. This type of regularization is a special case of the familiar Tikhonov regularization. There are other potential regularizations and problem formulations.[11] For example, Halperin and Tomecek (2008) propose a formulation as an entropy minimization problem. Here, the thinning matrix $\bar{M}$ is chosen relative to a prior matrix that reflects the risk ordering of the constituents, or other criteria.

The valuation formulae and the optimization are implemented in Matlab. The computation of (47) and (48) is based on the Fourier method described in §6.4. We use the Mosek Matlab toolbox to address (49). The solution to the matrix calibration problem is found within seconds. All single-name market swap rates are matched perfectly on each of the 21 trading days of September 2008.

The fitted matrix $\bar{M}$ provides interesting information: an element $\bar{q}^{kn}$ of $\bar{M}$ represents the market-implied probability of name k to be the nth defaulter. Figure 2 shows the fitted probability that the real estate finance company, Residential Capital, is the first defaulter for each trading day in September 2008. To understand the time series behavior, we need to consider the default risk of Residential Capital relative to that of the other names in the index. According to market swap spreads, Residential Capital is the riskiest name in the CDX High Yield index for the first 18 days in the sample. The distances to the second and third riskiest names, media company Tribune Company and pulp and paper company Abitibi Inc., respectively, are large for the first 16 days. The gap to Abitibi narrows very quickly over the subsequent two days, with Abitibi taking over Residential Capital as the riskiest name for the last three trading days of September. With Abitibi becoming more risky than Residential Capital, Residential's probability of defaulting first decreases sharply.

Figure 2. Market-implied probability of Residential Capital LLC, a real estate finance company, to be the first defaulter in the CDX High Yield portfolio, for each of the 21 trading days of September 2008.



9. Estimating Hedge Sensitivities

A portfolio derivative investor is exposed to fluctuations in the position's value due to changes in the constituent's

risks. For example, a tranche protection seller is exposed to an increase in the default risk of any constituent. To hedge this exposure, we require the sensitivity of the position value to small changes of the value of a constituent single-name swap. The product of this sensitivity and the contract notional is the amount of single-name swap protection that must be bought or sold to hedge the exposure due to small changes in the market spread of the constituent.[12] For index or tranche j, we estimate the sensitivity with respect to name k by the Delta

$$\Delta^{jk} = \frac{\Delta C^{jk}}{\Delta C^k}, \tag{50}$$

where ΔC^{jk} is the change in the value of tranche j due to a small, fixed shift in the swap rate of firm k, and ΔC^k is the corresponding change in the value of the single-name swap. From the perspective of the protection seller, the value of a single-name, index, or tranche swap is equal to the value of the future premium payments minus the value of the potential protection payments; see §7. At contract inception, the value is zero. We use the following scheme to estimate the change in value:

(1) Calibrate λ and the $\bar{Z}^k$ by fitting $\theta = (\lambda_0, \kappa, c, \delta)$ and $\bar{M} = (\bar{q}^{kn})$ from a set of market rates of single-name, index and tranche swaps as in §8.

(2) Shift the single-name swap rate $\mathrm{Mid}(k)$ of firm k by 0.1%.[13] Leave all other market rates unchanged. Replace θ by $\theta^\epsilon = (\lambda_0, \kappa, \epsilon c, \delta)$ for a scaling factor $\epsilon > 0$; all other parameters are set to their values calibrated in Step (1). Holding $\bar{M}$ fixed, recalibrate λ by choosing ε to minimize the single-name swap fitting error:

$$\min_{\epsilon>0} \sum_{k=1}^{m} \left(\sum_{n=1}^{m} \bar{q}^{kn} (V^n_{\theta^\epsilon} \mathrm{Mid}(k) + X^n_{\theta^\epsilon} l) - W \mathrm{Mid}(k) \right)^2 + \sum_{k=1}^{m} \sum_{n=1}^{m} (\bar{q}^{kn})^2. \tag{51}$$

(3) For each of the single-name, index, and tranche swaps, calculate the value of the future premium and protection payments based on $(\theta^\epsilon, \bar{M})$ and the pre-shift market rates as in §7. For swap k, the difference between the values is ΔC^k. For an index or tranche swap j, the difference between the values is ΔC^{jk}.

The estimation procedure assumes that the thinning matrix $\bar{M}$ remains unchanged when the spread of a constituent name is shifted. The matrix $\bar{M}$ describes the relative riskiness of the constituents, and we argue that a small shift in the spread of a (typically risky) name does not perturb the default ordering of names. Rather, a change in the risk of a constituent has an impact on the level of portfolio credit risk, which is governed by the portfolio intensity λ. We measure this impact through the intensity parameter c, which represents the base default rate in the portfolio. This choice is economically meaningful, computationally convenient, and leads to well-performing hedges, as we demonstrate below.[14] The recalibration Step (2) determines the proportional adjustment ϵ to the intensity parameter c calibrated in Step (1) that is necessary to fit the set of shifted single-name spreads. The re-pricing Step (3) determines the change in the index and tranche values associated with the adjustment ϵ to c obtained in Step (2). It exploits the sensitivity of index and tranche rates to ϵ, which increases with tranche seniority.

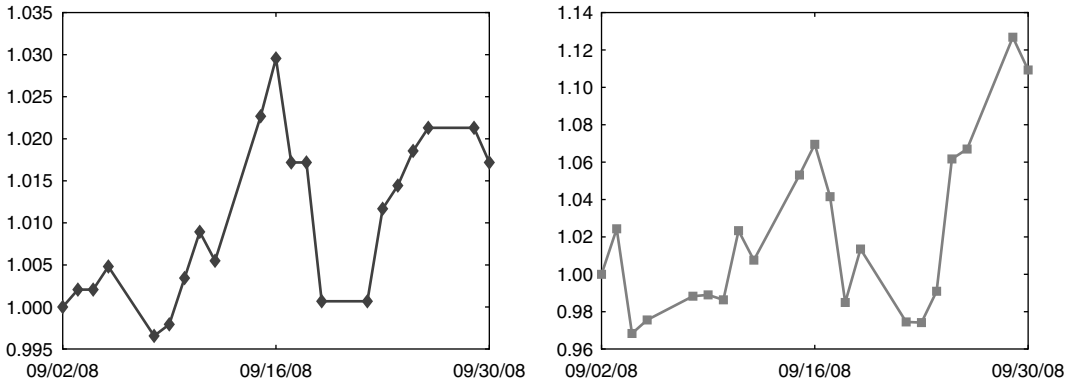
We evaluate the single-name hedges estimated by this scheme based on the market data for September 2008. We consider a protection selling position in a $T = 5$ year equity tranche referenced on the CDX High Yield portfolio with notional equal to one unit. For each of the 21 trading days of September 2008, we estimate the sensitivities (50) for Tribune Company. Media company Tribune is among the riskiest names in the CDX High Yield portfolio and therefore is a natural candidate for hedging.

Figure 3 shows the normalized tranche and Tribune swap prices. The time-series behavior reflects the default of Lehman on the 15th and other events. The goal is to offset the fluctuations of the tranche value due to the volatility of Tribune spreads by entering into a protection buying position in the single-name swap referenced on Tribune. The Tribune position is set up on September 2, 2008 and is re-balanced on each consecutive day according to the estimated hedge sensitivities. Because the tranche notional is equal to one, the sensitivity represents the notional of single-name protection to be bought on each day.

To measure the performance of the hedges, we consider the daily profit or loss (P/L) of the hedged tranche position, which is the daily change in the value of the hedged tranche position. The P/L measures how well the single-name position counteracts market fluctuations of the tranche value due to changes of Tribune spreads. We prefer hedges that generate low P/L volatility as measured by the standard deviation of the P/L.

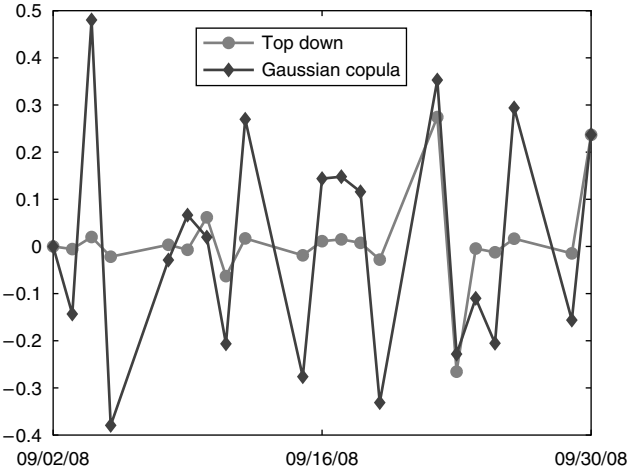
To get a sense for the relative performance of the hedges, we contrast the P/L generated by the top-down approach with that generated by the Gaussian copula model, which is the financial industry standard for tranche pricing and hedging.[15] The sensitivities (50) generated by the industry implementation of the copula model are obtained from Bloomberg (function CDST). They are based on the same single-name swap rate data that we used for the estimation of the sensitivities (50) in our top-down model setting. Figure 4 shows the time series of the daily incremental P/L generated by the two approaches. The graph indicates that the top-down hedges generated by our method outperform the copula based hedges. The top-down hedges lead to lower P/L volatility (standard deviation of 0.102) than the copula hedges (standard deviation of 0.240).[16] In other words, the top-down hedges offset the fluctuations of tranche rates due to the volatility of Tribune spreads better than do the copula hedges. The top-down hedges perform particularly well around Lehman Brother's default on September 15, 2008, when tranche rates fluctuate significantly (Figure 3). We conclude that the top-down

Figure 3. Market swap prices, normalized by their value on 9/2/2008.



Note. Left panel: Equity tranche upfront rate. *Right panel:* Swap spread of Tribune Company.

Figure 4. Daily incremental P/L of hedged tranche during September 2008.



approach enables better single-name hedging than the copula approach.

10. Conclusion

Multiname credit derivatives play an important role in credit markets: they provide tailored insurance against losses in a reference portfolio of defaultable assets and therefore allow investors to mitigate their exposure to correlated default risk. A popular approach to analyzing these derivatives is to specify the dynamics of 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 computationally tractable derivative pricing problem in which the underlying follows a point process with given intensity.

This paper develops a top-down approach that extends the reach of any stand-alone portfolio point process model to the constituents. This approach allows the researcher to consistently analyze securities that are referenced on the portfolio or on individual constituent firms. It facilitates important applications that cannot be addressed with a stand-alone portfolio point process model—for example, constituent risk hedging. At the center of the top-down approach is random thinning, which allocates portfolio-level risk to the constituents. The allocation is specified by a thinning process, which exists uniquely for any portfolio point process model and is a probabilistic model of the next-to-default. It gives rise to a simple formula for the constituent default probability, which is the key to single-name applications such as hedging. Tractable transform methods are developed to evaluate this formula for many models used in the literature.

Extensive market calibration experiments document the computational feasibility and the utility of the top-down approach. An empirical study demonstrates the superior performance of the constituent name hedges generated by the top-down approach for portfolio derivatives referenced on the CDX High Yield portfolio during September 2008, a month that witnessed significant price volatility due to the default of Lehman Brothers and the collapse of AIG. The hedges generated by the top-down approach outperform the hedges generated by the Gaussian copula model, which is the standard model for portfolio derivatives pricing and hedging used in the financial industry.

Appendix A. Proofs

Proof of Lemma 2.1. The process $A = \int_0^{\cdot} \lambda_s \, ds$ defined by formula (4) starts at zero and is nondecreasing. It is adapted and has continuous sample paths almost surely-$\mathbb{P}$ and is therefore predictable. Using formula (1) the process $N - A$ is the sum of martingales $N^k - A^k$ and thus a martingale itself. By the uniqueness of the Doob-Meyer decomposition, A is the compensator to N. □

Proof of Proposition 3.1 We prove a stronger result that only requires that a single-name compensator A^k has continuous paths almost surely-$\mathbb{P}$.

The random measure $dA_t^k(\omega)$ on $\mathbb{R}_+$ induced by A^k is absolutely continuous with respect to the random measure $dA_t(\omega)$ on $\mathbb{R}_+$ induced by A, almost surely-$\mathbb{P}$. This is

because $dA_t(\omega)$ is equal to the sum over k of the positive measures $dA_t^k(\omega)$ almost surely-$\mathbb{P}$, using the properties of the Doob-Meyer decomposition in analogy to Lemma 2.1. The corresponding Radon-Nikodym derivative process dA^k/dA is predictable thanks to the predictability of the processes A^k and A; see Dellacherie and Meyer (1982, Chapter VI, Theorem 68). It is also unique. We are interested in the representation of the process $Z^k = dA^k/dA$ in terms of a concrete limit procedure. We apply the results in §4 of Airault and Föllmer (1974), taking the nonnegative supermartingales X and Y considered there to be $1 - N^k$ and $m - N$, respectively. To this end, note that both supermartingales are regular potentials of class (D). Their continuous compensators are A^k and A, respectively.

We consider the product space $\Omega^* = \Omega \times (0, \infty]$, the product sigma-field $\mathcal{F}^*$ and the sigma-field $\mathcal{P}$ of predictable sets in Ω^*. We introduce the measures on $\mathcal{P}$ associated with the supermartingales at hand: for $B \in \mathcal{F}_t$ and $t \geqslant 0$,

$$\mu^k(B \times (t, \infty]) = \mathbb{E}((1 - N_t^k)1_B) = \mathbb{E}((H_t^k - A_t^k)1_B), \quad \text{(A1)}$$

$$\mu(B \times (t, \infty]) = \mathbb{E}((m - N_t)1_B) = \mathbb{E}((H_t - A_t)1_B), \quad \text{(A2)}$$

where H^k resp. H is the martingale in the Doob-Meyer decomposition of N^k resp. N. Dellacherie and Meyer (1982, p. 86). associate to A the measure ν on $\mathcal{F}^*$ given by

$$\nu(U) = E\left(\int_0^\infty U_s \, dA_s\right) \quad \text{(A3)}$$

for nonnegative and adapted processes U. We remark that μ and ν coincide on $\mathcal{P}$ since they coincide on the generating sets $B \times (t, \infty]$ with $B \in \mathcal{F}_t$. Thus, μ is simply the restriction of ν to $\mathcal{P}$.

Because $1 - N^k$ is a potential of class (D) and $dA_t^k(\omega)$ is absolutely continuous with respect to $dA_t(\omega)$ almost surely-$\mathbb{P}$, Lemma 4.1 in Airault and Föllmer (1974) implies that μ^k is absolutely continuous with respect to μ. The corresponding density is predictable as the density of two measures defined on $\mathcal{P}$.

Define the infinitesimal operator Z^k of the supermartingale $(m - N)$, which associates to the supermartingale $(1 - N^k)$ the value

$$Z_t^k(\omega) = \lim_{\epsilon \to 0} \frac{\mathbb{E}(N_{t+\epsilon}^k - N_t^k \mid \mathcal{F}_t)}{\mathbb{E}(N_{t+\epsilon} - N_t \mid \mathcal{F}_t)}(\omega) \quad \text{(A4)}$$

at points $(\omega, t) \in \Omega^*$ where the limit exists. Then, by Theorems 4.2 and 4.7 in Airault and Föllmer (1974), the infinitesimal operator Z^k exists and

$$\frac{dA^k}{dA} = \frac{d\mu^k}{d\mu} = Z^k \qquad \mu\text{-a.s.}, \quad \text{(A5)}$$

where the densities $dA^k/dA(\omega, \cdot)$ are chosen such that the process dA^k/dA is predictable.

If there are progressively measurable processes λ^k such that (3) holds almost surely-$\mathbb{P}$, then (A5) implies that

$$\lambda^k = \lambda Z^k \qquad \mu\text{-a.s.}, \quad \text{(A6)}$$

where $\lambda = \lambda^1 + \cdots + \lambda^m$ satisfies (2) almost surely-$\mathbb{P}$. This completes the proof. □

Proof of Proposition 4.1 This is a consequence of Corollary 4.11 in Airault and Föllmer (1974), Fubini's theorem and Proposition 3.1. Alternatively, we can directly exploit the martingale property of $N^k - \int_0^{\cdot} \lambda_s^k \, ds$ along with Fubini's theorem and Proposition 3.1. □

Proof of Proposition 4.3 From Proposition 4.1 we obtain

$$\mathbb{P}(t < \tau^k \leqslant T \mid \mathcal{F}_t) = \int_t^T \sum_{n=N_t+1}^m \mathbb{E}(\lambda_s M_s^{kn} 1_{\{N_s = n-1\}} \mid \mathcal{F}_t) \, ds.$$

It remains to calculate the conditional expectation in the integrand. We have

$$\begin{aligned}
&\mathbb{E}(\lambda_s M_s^{kn} 1_{\{N_s = n-1\}} \mid \mathcal{F}_t) \\
&\quad = -\partial_z \mathbb{E}(\exp(-z\lambda_s M_s^{kn}) 1_{\{N_s = n-1\}} \mid \mathcal{F}_t)|_{z=0} \\
&\quad = -\partial_z \varphi_t(n - 1 - N_t, z, s, \lambda_s M_s^{kn})|_{z=0},
\end{aligned}$$

where the second line follows from Equation (25). □

Appendix B. Extension of Results

This appendix discusses the extension of the results in §§3 and 4. We consider a right-continuous and complete subfiltration $\bar{\mathbb{F}} = (\bar{\mathcal{F}}_t)_{t \geqslant 0}$ of $\mathbb{F}$ relative to which the portfolio default counting process N is adapted but the constituent default times τ^k are not necessarily stopping times. To apply the results of Airault and Föllmer (1974) in this setting, we take the supermartingale X considered by Airault and Föllmer (1974) to be the optional projection of the $\mathbb{F}$-supermartingale $(1 - N^k)$ onto $\bar{\mathbb{F}}$. That X is an $\bar{\mathbb{F}}$-supermartingale follows from iterated expectations; see the lemma in Protter (2004, p. 368). Furthermore, we take the supermartingale Y considered by Airault and Föllmer (1974) to be $m - N$, which is an $\bar{\mathbb{F}}$-supermartingale because N is $\bar{\mathbb{F}}$-adapted. Note that both X and Y are $\bar{\mathbb{F}}$-potentials of class (D). The $\bar{\mathbb{F}}$-compensator $\bar{A}^k$ of X is given by the dual predictable projection of the $\mathbb{F}$-compensator A^k onto $\bar{\mathbb{F}}$, and the process Y has $\bar{\mathbb{F}}$-compensator $\bar{A} = \bar{A}^1 + \cdots + \bar{A}^m$ because Y is equal to the sum over k of the optional projection of $(1 - N^k)$. If A^k is absolutely continuous with respect to Lebesgue measure, then so is $\bar{A}^k$.

Define the measures μ^k and μ in analogy to (A1) and (A2) for $B \in \bar{\mathcal{F}}_t$. Because the random measure $d\bar{A}_t^k(\omega)$ induced by $\bar{A}^k$ is absolutely continuous with respect to the random measure $d\bar{A}_t(\omega)$ induced by $\bar{A}$ almost surely-$\mathbb{P}$,

Lemma 4.1 in Airault and Föllmer (1974) implies that μ^k is absolutely continuous with respect to μ. Let

$$\bar{Z}_t^k(\omega) = \lim_{\epsilon\to 0} \frac{\mathbb{E}(X_t - X_{t+\epsilon} \mid \bar{\mathcal{F}}_t)}{\mathbb{E}(Y_t - Y_{t+\epsilon} \mid \bar{\mathcal{F}}_t)}(\omega) = \lim_{\epsilon\to 0} \frac{\mathbb{E}(N_{t+\epsilon}^k - N_t^k \mid \bar{\mathcal{F}}_t)}{\mathbb{E}(N_{t+\epsilon} - N_t \mid \bar{\mathcal{F}}_t)}(\omega) \tag{B1}$$

at points $(\omega, t) \in \Omega^*$ where the limit exists. The second equality in (B1) uses the fact that $\bar{N}^k$ is the optional projection of N^k onto $\bar{\mathbb{F}}$. Then by Theorems 4.2 and 4.7 in Airault and Föllmer (1974), the infinitesimal operator $\bar{Z}^k$ exists and

$$\frac{d\bar{A}^k}{d\bar{A}} = \frac{d\mu^k}{d\mu} = \bar{Z}^k \qquad \mu\text{-a.s.} \tag{B2}$$

The process $\bar{Z}^k$ is the conditional expectation under μ of the density Z^k in (A5). This amounts to taking $\bar{Z}^k$ as the predictable projection of Z^k onto $\bar{\mathbb{F}}$.

To prove the default probability formula (17) relative to $\bar{\mathbb{F}}$, we proceed as in the proof of Proposition 4.1: we apply Corollary 4.11 in Airault and Föllmer (1974) relative to $\bar{\mathbb{F}}$, and then use Fubini's theorem and formula (B2).

To prove formula (20), we mimic the proof of Proposition 4.3: we start with formula (17), apply formula (19), note that $\bar{\lambda} = \lambda$ because λ is assumed to be $\bar{\mathbb{F}}$-adapted in the context of (20), set $t = 0$, and then use the definition of φ.

Appendix C. Optimization Problem

This appendix discusses the optimization problem (45). We start by considering the unregularized problem, given by

$$\min_{(\bar{q}^{kn})} \sum_{k=1}^m \left(\left(\sum_{n=1}^m \bar{q}^{kn} a^{kn} \right) - b^k \right)^2,$$

$$\text{such that} \quad \sum_{k=1}^m \bar{q}^{kn} = 1, \quad n = 1, \ldots, m$$

$$\sum_{n=1}^m \bar{q}^{kn} = 1, \quad k = 1, \ldots, m$$

$$\bar{q}^{kn} \geqslant 0, \quad k, n = 1, \ldots, m,$$

where $a^{kn} = V_{\theta^*}^n \mathrm{Mid}(k) + X_{\theta^*}^n l$ and $b^k = W\mathrm{Mid}(k)$. Letting $x^{kn} = \bar{q}^{kn} a^{kn}$ and noting that $a^{kn} > 0$, this problem can be rewritten as

$$\min_{(x^{kn})} \sum_{k=1}^m \left(\sum_{n=1}^m x^{kn} \right)^2 - 2 \sum_{k=1}^m b^k \sum_{n=1}^m x^{kn},$$

$$\text{such that} \quad \sum_{k=1}^m \frac{x^{kn}}{a^{kn}} = 1, \quad n = 1, \ldots, m$$

$$\sum_{n=1}^m \frac{x^{kn}}{a^{kn}} = 1, \quad k = 1, \ldots, m$$

$$x^{kn} \geqslant 0, \quad k, n = 1, \ldots, m.$$

Now stack the matrix (x^{kn}) column-wise into an m^2-vector x. Then the problem can be written as

$$\min_x \ \frac{1}{2} x^\top Q x + b^\top x, \tag{C1}$$

$$\text{such that} \quad Cx = e, \qquad Dx = e, \quad x \geqslant 0, \tag{C2}$$

where e is an m-vector of ones and Q is an $(m^2 \times m^2)$ block diagonal matrix in which the diagonal entries are given by an $(m \times m)$ matrix of ones, denoted E. Furthermore, b is an m^2-vector given by $b = (b^1, b^1, \ldots, b^1, b^2, b^2, \ldots, b^2, \ldots, b^m, b^m, \ldots, b^m)^\top$ and C is a sparse $(m \times m^2)$ matrix with all elements equal to zero except the $(i, (j-1)m + i)$th element, which is equal to $1/a^{ji}$, for $j = 1, \ldots, m$. Similarly, D is a sparse $(m \times m^2)$ matrix with all elements equal to zero except the $(i, (i-1)m + j)$th element, which is equal to $1/a^{ji}$, for $j = 1, \ldots, m$. The problem (C1)–(C2) is a standard convex programming problem and, more specifically, a constrained quadratic programming problem. The special structure of the matrix Q implies that it is rank deficient and not positive definite. Therefore, the problem (C1)–(C2) does not have a unique solution.

When we add the regularization term to (45) and rewrite the regularized problem as a convex program as above, we see that the diagonal entries of the $(m^2 \times m^2)$ matrix Q in the resulting convex program are given by $E + I$, where I is the $(m \times m)$ identity matrix. Thus, Q is now of full rank and positive definite, so the resulting convex program has a unique solution.

Endnotes

1. We refer to Protter (2004) for technical definitions not offered here.
2. See, for example, Proposition 14.1.VI in Daley and Vere-Jones (2008) in this regard.
3. The statement that for all t, the limit (5) exists $\mathbb{P}$-almost surely is inadequate to our needs since it is consistent with $Z_t^k(\omega)$ being ill defined on a set of positive μ-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.
4. For certain filtrations $\mathbb{F}$, the representation (13) is necessary and sufficient for Z^k to be predictable; see Jacod (1975). This includes the case where $\mathbb{F}$ is the filtration generated by N.
5. Unless stated otherwise, in this section measurability properties refer to the filtration $\bar{\mathbb{F}}$.
6. See Dellacherie and Meyer (1982, pp. 43–44) for further details.
7. The optional projection of a default indicator process onto a subfiltration appears also in other contexts. For example, it plays a key role in the analysis of single-name credit risk under incomplete information; see Elliott et al. (2000) and Giesecke (2006).

8. See Dellacherie and Meyer (1982, pp. 43–44) for further details.
9. In market practice, the loss at default is assumed to be constant, with a typical value of 60%.
10. We follow the market practice of adjusting the swap rates to account for differences in the terms of the index and constituent swap contracts. For example, a credit event that triggers a payoff in a single-name credit swap need not lead to a payoff in an 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. 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 $\mathrm{Mid}(k)$ is the arithmetic average of the adjusted market bid and ask quotes for the swap referenced on firm k.
11. We have experimented with several alternative regularizations for the problem (49). The hedge ratios estimated in §9 below, which are based on the solution to the problem (49), were found to be insensitive to the regularization used.
12. Other sensitivities can be considered, including the sensitivity of the value of a tranche position to changes in the index swap rate. Estimating this sensitivity does not require random thinning and can be performed with a stand-alone model of portfolio loss. See Cont and Kan (2009).
13. We have experimented with different amounts including 0.01% and 1% but found the estimated hedge ratios (50) to be insensitive to the magnitude of the shift used.
14. Alternative choices might lead to different hedges. The dependence of hedges on the scheme used to estimate them is a property that applies to all portfolio derivatives valuation models, including bottom-up formulations. See Andersen and Sidenius (2005), Lopatin (2008), and Halperin and Tomecek (2008) for a discussion. Thus, for a given valuation model the hedges are not unique.
15. The copula model goes back to Li (2000). We use the copula model as a benchmark for comparison purposes only; we do not endorse its use for pricing or hedging.
16. The mean of the P/L generated by the top-down approach is 0.010, while the mean generated by the copula model is 0.003. Thus, the hedged position generated by the top-down approach is more profitable, on average. We do not expect the mean to be 0 because we do not hedge every portfolio constituent.

Acknowledgments

This is a revision of an earlier paper by the first two authors. An extended version of the valuation models from that earlier paper now appears as Errais et al. (2010). The authors thank Greg Anderson, Aaron Brown, Eymen Errais, Steve Evans, Bjorn Flesaker, Hans Föllmer, Igor Halperin, Steven Hutt, Jyh-Huei Lee, Michael Saunders, Thorsten Schmidt, and Pascal Tomecek for illuminating discussions. They also thank seminar participants for their comments at the 2007 Applied Probability Society Conference, the Fields Institute for Research in Mathematical Sciences, 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 2005 Newton Institute Conference on Developments in Quantitative Finance, the 2006 Numerical Methods for Finance Conference, and the 2005 Quantitative Methods in Finance Conference. They are deeply grateful to Stefan Weber for sharing his insights, as well as to two anonymous referees and the associate editor for their constructive comments.

References

Airault, H., H. Föllmer. 1974. Relative densities of semimartingales. *Inventiones Math.* **27** 299–327.
Andersen, L., J. Sidenius. 2005. CDO pricing with factor models: Survey and comments. *J. Credit Risk* **1**(3) 71–88.
Arnsdorf, M., I. Halperin. 2008. BSLP: Markovian bivariate spread-loss model for portfolio credit derivatives. *J. Comput. Finance* **12** 77–100.
Brigo, D., A. Pallavicini, R. Torresetti. 2007. Calibration of CDO tranches with the dynamical generalized-Poisson loss model. *Risk* **20**(June, 5) 70–75.
Chen, Z., P. Glasserman. 2008. Sensitivity estimates for portfolio credit derivatives using Monte Carlo. *Finance and Stochastics* **12** 507–540.
Collin-Dufresne, P., R. Goldstein, J. Hugonnier. 2004. A general formula for the valuation of defaultable securities. *Econometrica* **72** 1377–1407.
Cont, R., Y. H. Kan. 2009. Dynamic hedging of portfolio credit derivatives. Working paper, Columbia University, New York.
Cont, R., A. Minca. 2008. Extracting portfolio default rates from CDO spreads. Working paper, Columbia University, New York.
Daley, D., D. Vere-Jones. 2008. *An Introduction to the Theory of Point Processes*, Vol. 2. Springer-Verlag, New York.
Das, S., D. Duffie, N. Kapadia, L. Saita. 2007. Common failings: How corporate defaults are correlated. *J. Finance* **62** 93–117.
Davis, M., V. Lo. 2001. Modeling default correlation in bond portfolios. C. Alexander, ed. *Mastering Risk Volume 2: Applications*. Prentice Hall, Englewood Cliffs, NJ, 141–151.
Dellacherie, C., P.-A. Meyer. 1982. *Probabilities and Potential*. North Holland, Amsterdam.
Ding, X., K. Giesecke, P. Tomecek. 2009. Time-changed birth processes and multiname credit derivatives. *Oper. Res.* **57**(4) 990–1005.
Duffie, D., N. Garleanu. 2001. Risk and valuation of collateralized debt obligations. *Financial Analysts J.* **57**(1) 41–59.
Duffie, D., J. Pan, K. Singleton. 2000. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* **68** 1343–1376.
Duffie, D., M. Schroder, C. Skiadas. 1996. Recursive valuation of defaultable securities and the timing of resolution of uncertainty. *Ann. Appl. Probab.* **6** 1075–1090.
Eckner, A. 2009. Computational techniques for basic affine models of portfolio credit risk. *J. Comput. Finance* **15** 63–97.
Elliott, R., M. Jeanblanc, M. Yor. 2000. On models of default risk. *Math. Finance* **10**(2) 179–195.
Errais, E., K. Giesecke, L. Goldberg. 2010. Affine point processes and portfolio credit risk. *SIAM J. Financial Math.* **1** 642–665.
Giesecke, K. 2006. Default and information. *J. Econom. Dynam. Control* **30**(11) 2281–2303.
Giesecke, K., S. Zhu. 2007. The correlation-neutral measure for portfolio credit. Working paper, Stanford University, Stanford, CA.

Halperin, I., P. Tomecek. 2008. Climbing down from the top: Single name dynamics in top-down credit models. Working paper, JP Morgan, New York.

Jacod, J. 1975. Multivariate point processes: Predictable projection, radon-nikodym derivatives, representation of martingales. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* **31** 235–253.

Jarrow, R. A., F. Yu. 2001. Counterparty risk and the pricing of defaultable securities. *J. Finance* **56**(5) 555–576.

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

Lando, D. 1998. On Cox processes and credit risky securities. *Rev. Derivatives Res.* **2**(2–3) 99–120.

Leippold, M., L. Wu. 2002. Asset pricing under the quadratic class. *J. Financial Quant. Anal.* **37** 271–295.

Li, D. X. 2000. On default correlation: A copula function approach. *J. Fixed Income* **9** 43–54.

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

Lopatin, A. 2008. A simple dynamic model for pricing and hedging heterogeneous CDOs. Working paper, Numerix, Warrenville, IL.

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

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

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

Protter, P. 2004. *Stochastic Integration and Differential Equations*. Springer-Verlag, New York.

Tavella, D., M. Krekel. 2006. Pricing *n*th-to-default credit derivatives in the PDE framework. Working paper, Octanti Associates, San Francisco.

Zhou, R. 2009. A multi-portfolio model for bespoke CDO pricing. Working paper, Depository Trust & Clearing Corporation.