PORTFOLIO CREDIT RISK:
TOP DOWN VS. BOTTOM UP APPROACHES

Kay Giesecke*

Stanford University

February 8, 2008†

Published in

Frontiers in Quantitative Finance: Credit Risk and Volatility Modeling

Edited by R. Cont, Wiley 2008

Abstract

Dynamic reduced form models of portfolio credit risk can be distinguished by the way in which the intensity of the default process is specified. In a bottom up model, the portfolio intensity is an aggregate of the constituent intensities. In a top down model, the portfolio intensity is specified without reference to the constituents. This expository article contrasts these modeling approaches. It emphasizes the role of the information filtration as a modeling tool.

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

†I would like to thank Rama Cont, Lisa Goldberg and Jack Kim for helpful discussions, comments and suggestions.

1 Introduction

A model of portfolio credit risk has three elements: a filtration that represents the observable information, a default process that counts events in the portfolio and a distribution for the financial loss at an event. Reduced form models of portfolio credit risk can be distinguished by the way in which the intensity of the default process is specified. In a *bottom up* model, the portfolio intensity is an aggregate of the constituent intensities. In a *top down* model, the portfolio intensity is specified without reference to the constituents. The constituent intensities are recovered by random thinning. This expository article contrasts the two modeling approaches.

2 Portfolio credit models

Consider a portfolio of credit sensitive securities such as loans, bonds or credit swaps. The ordered portfolio default times are represented by a sequence of stopping times $T^n > 0$ that is strictly increasing to infinity and defined on a complete probability space $(\Omega, \mathcal{F}, P)$ with a right-continuous and complete filtration $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ that represents the information flow. The random variable T^n represents the nth default time in the portfolio. Depending on the context, P can be the actual probability or a risk-neutral measure. Let N be the process that counts default events, given by

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

A portfolio credit model is a specification of the filtration $\mathbb{F}$, default process N and distribution for the loss at an event. For a given filtration, the default process N is specified in terms of its compensator, which is the non decreasing predictable process A such that $N - A$ is a local martingale. The compensator embodies the expected upward tendency of the default process. Meyer (1966) shows that in the limit,

$$A_t = \lim_{\epsilon \downarrow 0} \frac{1}{\epsilon} \int_0^t E\big[N_{s+\epsilon} - N_s \,|\, \mathcal{F}_s\big] ds, \tag{2}$$

weakly in L^1. Formula (2) emphasizes the dependence of the compensator on the filtration. The filtration, the probabilistic properties of the default times, and the analytic properties of the compensator are closely related. If the times are predictable, i.e. if an event is announced by a sequence of pre-default times, then A is equal to N. As an example, consider the familiar first passage credit models that descend from Black & Cox (1976). Here, a firm defaults if its continuous firm value process falls below a constant barrier. This definition of the default event generates a predictable default time. If, as in Duffie & Lando (2001) or Giesecke (2006), the available information is insufficient to determine the precise value of the firm's assets or default barrier, then the default times are totally inaccessible or unpredictable. In this case defaults come without warning and the compensator A is continuous. Unpredictable default times can conveniently be specified in terms of a non negative, adapted *intensity* λ that satisfies

$$A_t = \int_0^t \lambda_s \, ds, \tag{3}$$

almost surely. Together, formulae (2) and (3) show that the intensity is the conditional portfolio default rate in the sense that $\lambda_t \Delta$ is approximately equal to $P[N_{t+\Delta} - N_t = 1 \,|\, \mathcal{F}_t]$ for small Δ. If the compensator is of the form (3), then the portfolio credit model is intensity based. In a top down model the process λ is specified directly. In a bottom up model, λ is an aggregate of constituent intensity processes.

3 Information and specification

The structure of the information filtration $\mathbb{F}$ determines the key properties of a portfolio credit model. The filtration must always be fine enough to distinguish the arrival of events. Therefore the smallest filtration that supports a portfolio credit model is the filtration generated by the default process N itself. Bottom up and top down model specifications are based on distinct filtrations, which explains many of the structural differences between them.

3.1 Bottom up models

A bottom up model filtration usually contains much more information than the minimal filtration. It is always fine enough to distinguish the identity of each defaulter so that the constituent default times τ^k are stopping times. The filtration may contain additional information about the prices of single- and multi-name derivatives, macro-economic variables and other systematic and idiosyncratic risk factors.

A constituent default time τ^k generates a default process N^k that is zero before default and one afterwards. If the portfolio default process $N = \sum_k N^k$ is intensity based, then so is each constituent default process. In this case, there is a strictly positive intensity process λ^k that represents firm k's conditional default rate in the sense that $N^k - \int_0^\cdot (1 - N_s^k)\lambda_s^k ds$ is a martingale.

The researcher specifies the model filtration $\mathbb{F}$ and the constituent intensity processes λ^k. The dependence structure of all firms must be build into each of the constituent intensity processes. Empirical observation suggests distinguishing two sources of firm dependence. First, firms are exposed to common or correlated economic factors such as interest rates or commodity prices. The variation of these factors generates correlated changes in firms' default rates, and the cyclical pattern in the time-series behavior of aggregate default rates. Second, due to the complex web of business, legal and informational relationships in the economy, defaults have a direct impact on the default rates of the surviving firms. For example, the collapse of automotive manufacturer Delphi in 2005 severely affected General Motors, whose production critically depended on Delphi's timely supply of parts. In response to the event, investors immediately demanded a higher default insurance premium for General Motors, reflecting the sudden increase in GM's likelihood to fail. Collin-Dufresne, Goldstein & Hugonnier (2004), Jorion & Zhang (2007*b*) and Jorion & Zhang (2007*a*) show that this episode is not an isolated case.

We illustrate several constituent intensity specifications. Each example specification incorporates different channels for default correlation.

Example 3.1. Let $\mathbb{F}$ be the filtration generated by the constituent default processes. For a deterministic function $c^k > 0$ that models the base intensity, set

$$\lambda^k = c^k + \sum_{j \neq k} \delta^{kj} N^j$$

see Giesecke & Weber (2006) and Kusuoka (1999). At each event, a term is added to the intensity that reflects the response of firm k's default rate to the event. The sensitivity of firm k to the default of firm j is modeled by the deterministic function $\delta^{kj} \geq 0$. If these sensitivities are zero, then the intensity varies only deterministically and the constituent default times are independent. □

Example 3.2. Let $\mathbb{F}$ be the filtration generated by the constituent default processes, a systematic risk factor X and a collection of idiosyncratic risk factors X^k that are independent of one another and independent of X. For a deterministic function α^k that describes the exposure of firm k to the factor X,

$$\lambda^k = \alpha^k X + X^k. \tag{4}$$

All firms are sensitive to the systematic factor X. Movements of X generate correlated changes in firms' intensities. If the risk factors evolve independently of the firm default processes as in Das, Duffie, Kapadia & Saita (2007), Duffie & Garleanu (2001), Duffie, Saita & Wang (2006), Eckner (2007), Feldhütter (2007) and Mortensen (2006), then the specification (4) generates a doubly stochastic model. Here, *conditional* on a path of the systematic factor, firms default independently of one another. The specification (4) can be extended to include multiple common factors that model sectoral, regional or other risks. Papageorgiou & Sircar (2007) partition firms into homogeneous groups or sectors, and take (4) as the common intensity of firms in a given sector k. The factor X induces correlation between sectors. The factor X^k models sector-specific risk. □

Example 3.3. Let $\mathbb{F}$ be the filtration generated by the constituent default processes, a systematic risk factor X and a collection of idiosyncratic risk factors X^k that are independent of one another and independent of X. Let

$$\lambda^k = \alpha^k X + X^k + \sum_{j \neq k} \delta^{kj} N^j,$$

see Collin-Dufresne et al. (2004), Frey & Backhaus (2004), Giesecke & Weber (2004), Jarrow & Yu (2001), Schönbucher & Schubert (2001) or Yu (2007). This specification incorporates the sensitivity of firms to a common risk factor, and event feedback through the terms that reflect the default status of the firms in the portfolio. The factor sensitivity is ignored in Example 3.1. The doubly stochastic Example 3.2 ignores event feedback. □

Example 3.4. Let $\mathbb{F}$ be the filtration generated by the constituent default processes, a systematic risk factor X and a collection of idiosyncratic risk factors X^k that are independent of one another and independent of X. Let U be another systematic risk factor that is not observable, i.e. not adapted to $\mathbb{F}$. This frailty factor must be projected onto $\mathbb{F}$ to obtain an observable process $\widehat{U}$ given by $\widehat{U}_t = E[U_t \,|\, \mathcal{F}_t]$. For a deterministic function δ^k that specifies the exposure of firm k to U,

$$\lambda^k = \alpha^k X + X^k + \delta^k \widehat{U},$$

see Collin-Dufresne, Goldstein & Helwege (2003), Duffie, Eckner, Horel & Saita (2006), Delloye, Fermanian & Sbai (2006), Giesecke & Goldberg (2004) and Schönbucher (2004). The projection $\widehat{U}$ is updated with observable information. In particular, $\widehat{U}$ is revised at events since the filtration contains information about firms' default status. Since events are unpredictable, the projection jumps at an event time. A jump corresponds to Bayesian updating of investors' beliefs about the distribution of the frailty factor U. This updating leads to intensity dynamics that are qualitatively similar to those in Example 3.3. □

The constituent intensities λ^k determine the portfolio intensity λ. Since the portfolio default process N is the sum over k of the constituent default processes N^k and defaults occur at distinct dates almost surely, the portfolio intensity is given by

$$\lambda = \sum_k (1 - N^k)\lambda^k, \tag{5}$$

see Giesecke & Goldberg (2005, Proposition 5.1). At each event, a term drops out of the sum. The portfolio intensity λ is zero after all firms have defaulted.

3.2 Top down models

In a top down model the researcher specifies the portfolio intensity λ without reference to the constituents. The name dependence structure is implicit in this specification. The goal is an intensity model that is more parsimonious than the bottom up portfolio intensity (5), which follows a complicated process that is driven by the constituent processes and depends on all single name parameters. This is achieved by choosing a model filtration $\mathbb{F}$ that is coarser than the bottom up model filtration. Typically, the top down model filtration is not fine enough to distinguish the identity of a defaulter. This means that an event arrival can be observed, but not the identity of the defaulted name.

Example 3.5. Let $\mathbb{F}$ be the filtration generated by the portfolio default process N and a risk factor $X > 0$ that evolves independently of N. Set

$$\lambda = X(m - N). \tag{6}$$

The risk factor X generates stochastic variation in the portfolio intensity between arrivals. It models the sensitivity of the portfolio constituents to a common economic factor. Conditional on a path of X, N is a time inhomogeneous death process. If X is a weighted sum of independent processes, then N is the generalized Poisson process of Brigo, Pallavicini & Torresetti (2006). □

This example illustrates how a relatively coarse filtration supports a parsimonious portfolio intensity specification. Instead of describing the constituent intensities, we focus on the inter-arrival intensity. In other words, we change the perspective from the firm default times τ^k to the ordered default times T^n. The example also illustrates the connection between a top down and bottom up specification. The top down portfolio intensity (6) coincides with the portfolio intensity (5) generated by an *exchangeable* bottom up model for which $\lambda^k = X$ for all constituents k.

In Example 3.5, the response of the portfolio intensity to events does not represent feedback but merely an adjustment that accounts for the fact that the set of potential defaulters is reduced at an event. To incorporate event feedback, we need to allow for a more flexible inter-arrival intensity specification.

Example 3.6. Let $\mathbb{F}$ be the filtration generated by the portfolio default process and a collection of risk factors X^n that vanish for $n \geq m$,

$$\lambda = X^n 1_{\{N=n\}}. \tag{7}$$

The intensity is revised at an event. Between events, the stochastic evolution of the intensity is governed by the processes X^n. If $X^n = X(m-n)$ for an independent common risk factor X, we obtain the doubly stochastic death process of Example 3.5. If $X^n = X\beta^n(m-n)$ for a deterministic function β^n, then N is the bivariate spread loss process of Arnsdorf & Halperin (2007). If $X^n = X(c + \delta n)$ for constants c and δ, then N is the time-changed birth process of Ding, Giesecke & Tomecek (2006). If $X^n = c + \delta n$ for deterministic functions c and δ, then N is the inhomogeneous birth process of Kim (2007). If X^n is deterministic then N is a time inhomogeneous Markov jump process as in Cont & Minca (2008) and Laurent, Cousin & Fermanian (2007). If for constants c, λ_0, κ and δ we set

$$X_t^n = c + (\lambda_0 - c)e^{-\kappa t} + \delta \sum_{j=1}^{n} e^{-\kappa(t - T^j)}$$

then N is the Hawkes process of Errais, Giesecke & Goldberg (2006). Davis & Lo (2001), Giesecke & Tomecek (2005), Longstaff & Rajan (2007) and Lopatin & Misirpashaev (2007) propose further specifications of the model (7). □

The change in perspective supports parsimonious specifications of the portfolio intensity λ, and as we illustrate below, it also leads to computational tractability of portfolio risk measurement and portfolio derivatives valuation. However, the top down approach calls for further steps if we require models for the constituent names. Constituent intensities are generated by *random thinning*, see Giesecke & Goldberg (2005). Random thinning disintegrates the portfolio intensity process λ into its constituent intensity processes. It provides the inverse to the intensity aggregation formula (5).

We must be careful about the notion of default intensity of a constituent. If, as in Examples 3.5 and 3.6, the top down model filtration $\mathbb{F}$ is not fine enough to distinguish the identity of a defaulter, then the constituent default processes N^k are not observable, i.e. adapted to $\mathbb{F}$. Therefore, we consider the projections $\widehat{N}^k$ onto $\mathbb{F}$. This is similar to the projection of the frailty factor onto the observation filtration in Example 3.4. Random thinning allocates the portfolio intensity λ to the intensities $\widehat{\lambda}^k$ of the constituent default process projections $\widehat{N}^k$. In Giesecke & Goldberg (2005), it is shown that for each portfolio intensity model, there exists a predictable thinning process Z^k such that

$$\widehat{\lambda}^k = Z^k \lambda. \tag{8}$$

The value Z^k_t is the conditional probability at time t that name k is the next defaulter given a default is imminent. Therefore, the sum of the Z^k_t over k must equal one unless all names in the portfolio are in default. If all constituents are in default the thinning processes vanish. As the following examples illustrate, the constituent intensities $\widehat{\lambda}^k$ inherit the properties of the portfolio intensity λ. In particular, they reflect the dependence structure of the ambient portfolio.

Example 3.7. Let $\mathbb{F}$ be the filtration generated by the portfolio default process N and an independent systematic risk factor $X > 0$. Let the portfolio intensity $\lambda = X(m - N)$ be as in Example 3.5. Consider the thinning process given by

$$Z^k_t = \frac{S^k}{\sum_{k=1}^m S^k} 1_{\{t \leq T^m\}}$$

where the S^k are the credit swap spreads of the constituent names observed at time 0 for a short, fixed maturity. The portfolio intensity is distributed according to the relative spread of names. A name whose spread is relatively wide compared with other names in the portfolio is attributed a relatively large share of the portfolio intensity. We have

$$\widehat{\lambda}^k_t = \frac{S^k}{\sum_{k=1}^m S^k} X_t(m - N_t) \tag{9}$$

so the exposure of firm k to the common factor is determined by the relative spread. However, the single name swap spreads implied by the constituent intensities (9) are not guaranteed to match observed spreads. Consider the alternative model

$$Z^k_t = s^k 1_{\{t \leq T^m\}},$$

where the s^k are non negative parameters such that $\sum_{k=1}^m s^k = 1$. Given a calibration of X from the tranche market, choose the parameters s^k such that the constituent intensities generate model credit swap spreads that are close to the market-observed credit swap spreads S^k. This calibration problem becomes well-posed if the constituent spreads are uniformly adjusted for the index basis, and the adjustment is calibrated along with the parameters s^k. □

In the previous example the thinning is static. Further, the portfolio and constituent intensities do not incorporate event feedback. The response of the intensities to an event merely represents an adjustment for the reduction in the set of potential defaulters.

Example 3.8. Let $\mathbb{F}$ be the filtration generated by the portfolio default process N and an independent systematic risk factor $X > 0$. For constants $c > 0$ and $\delta \geq 0$, let the portfolio intensity $\lambda = X(c+\delta N)1_{\{N<m\}}$, which generates the time-changed birth process of Ding et al. (2006), see Example 3.6. This specification incorporates the feedback of events. Letting $T^0 = 0$, consider the thinning process

$$Z_t^k = \sum_{n=1}^{m} z^{kn} 1_{\{T^{n-1}<t\leq T^n\}}$$

where $(z^{kn})_{k,n=1,2,\ldots,m}$ is a doubly stochastic matrix of non negative constants. The parameter z^{kn} represents the probability that firm k is the nth defaulter. With each event arrival, the portfolio intensity, thinning processes and the constituent intensities

$$\widehat{\lambda}^k = Z^k X(c+\delta N)$$

are revised. While the thinning is constant between events, the portfolio and constituent intensities fluctuate with the common factor X. The doubly stochastic thinning matrix is chosen so that the model implied single name swap spreads match observed spreads, see Ding et al. (2006) and Halperin (2007). □

4 Default distribution

A portfolio credit derivative is a contingent claim on the portfolio loss due to defaults. To calculate derivative prices and portfolio risk measures such as value at risk, we require the distribution of portfolio loss at multiple future horizons. Below we contrast the calculation of this distribution in bottom up and top down model specifications. To simplify the exposition, we assume that the loss at an event is constant. Therefore we can focus on the distribution of the default process N and its components N^k.

4.1 Bottom up models

In a bottom up model the default process $N = \sum_k N^k$ is the aggregate of the constituent default processes N^k. It is natural to consider the constituent default processes first.

Define $A^k_t = \int_0^t \lambda^k_s ds$, where λ^k is the intensity of firm k. If the variable $\exp(A^k_T)$ is square integrable for some horizon T, then for any time $t \leq T$ before default we have the conditional survival probability formula

$$P\big[\tau^k > T \,|\, \mathcal{F}_t\big] = E^*\big[e^{-(A^k_T - A^k_t)} \,|\, \mathcal{F}^*_t\big] \tag{10}$$

where the expectation on the right hand side is taken under an absolutely continuous probability measure P^* defined by the density $\exp(A^k_T)(1 - N^k_T)$, see Collin-Dufresne et al. (2004). The probability P^* puts zero mass on paths for which default occurs before T. The conditional expectation is taken with respect to the filtration $(\mathcal{F}^*_t)$, which is the completion of the reference filtration $\mathbb{F}$ by the P^*-null sets. Formula (10) applies to all bottom up constituent intensity specifications discussed in Section 3.1 above. The measure and filtration changes are redundant for the doubly stochastic Example 3.2. In this case, formula (10) simplifies to the classic formula derived by Lando (1998).

The conditional expectation on the right hand side of equation (10) is a familiar expression in ordinary term structure modeling. It is analogous to the price at t of a zero coupon bond with unit face value and maturity T, assuming the short term rate of interest is λ. The calculation of this price is well understood for a wide range of parametric short rate specifications, including affine and quadratic models. Formula (10) thus extends the analytical tractability offered by existing term structure model specifications to the constituent default process N^k.

Example 4.1. Let $\mathbb{F}$ be the filtration generated by the constituent default processes N^k, a systematic risk factor X and a collection of idiosyncratic risk factors X^k that are independent of one another and independent of X. The risk factors are independent of the N^k. For a constant α^k, set $\lambda^k = \alpha^k X + X^k$. From formula (10) we get

$$P\big[\tau^k > T\big] = E\big[e^{-\alpha^k \int_0^T X_s ds}\big] E\big[e^{-\int_0^T X^k_s ds}\big]. \tag{11}$$

The two expectations on the right hand side can be calculated explicitly if the risk factors follow affine jump diffusions or quadratic diffusions. In these cases, each expectation is an exponentially affine or quadratic function of the initial value of the risk factor. □

It is challenging to calculate the distribution of the portfolio default process N in the general bottom up setting. This is particularly true for intensity models with event feedback, where the calculations often rely on the special structure of the intensity parametrization. The calculations are most tractable in the doubly stochastic Example 3.2, which explains the popularity of this specification. Here we exploit the fact that conditional on the common risk factor, the N^k are independent.

Example 4.2. Consider the doubly stochastic setting of Example 4.1. Define the integrated common factor $Z_t = \int_0^t X_s ds$ and let

$$p_k(T, z) = P\big[\tau^k > T \,|\, Z_T = z\big] = e^{-\alpha^k z} E\big[e^{-\int_0^T X^k_s ds}\big]$$

be the conditional survival probability of firm k given a realization of the integrated common factor. The conditional probability generating function of the constituent default process is given by

$$E\big[v^{N_T^k} \mid Z_T = z\big] = p_k(T,z)(1-v) + v, \quad v \in \mathbb{R}.$$

By iterated expectations and conditional independence, the probability generating function of the portfolio default process N is

$$E\big[v^{N_T}\big] = \sum_{k=0}^{n} v^k P\big[N_T = k\big] = \int V(T,z,v) f_T(z) dz \tag{12}$$

where $f_T(z)$ is the density function of Z_T and

$$V(T,z,v) = \prod_{k=1}^{n} E\big[v^{N_T^k} \mid Z_T = z\big] = \prod_{k=1}^{n} \big(p_k(T,z)(1-v) + v\big)$$

Expanding the polynomial $V(T,z,v) = V_0(T,z) + vV_1(T,z) + \ldots + v^n V_n(T,z)$, from formula (12) we get the distribution of the portfolio default process:

$$P\big[N_T = k\big] = \int V_k(T,z) f_T(z) dz.$$

If the common risk factor follows an affine jump diffusion or quadratic diffusion, then the Laplace transform of Z_T is exponentially affine or quadratic in X_0, respectively, and the density $f_T(z)$ can be obtained by numerical transform inversion. Extensions of the single factor model for λ^k to include multiple common factors that model sectoral, regional or other risks are conceptually straightforward, but require multi-dimenional numerical transform inversion and integration which tend to be computationally very expensive. □

4.2 Top down models

In a top down model the distribution of the portfolio default process N can be calculated directly in terms of the portfolio intensity λ. Let $A_t = \int_0^t \lambda_s ds$ be the compensator to N. If the variable $\exp(A_T)$ is square integrable for some horizon T and Y is an integrable random payoff at T, then for real z, v and $t \leq T$ the default process transform

$$E\big[e^{izY + iv(N_T - N_t)} \mid \mathcal{F}_t\big] = E^v\big[e^{izY - (1-e^{iv})(A_T - A_t)} \mid \mathcal{F}_t\big] \tag{13}$$

where i is the imaginary unit and the expectation on the right hand side is taken under an equivalent complex measure P^v defined by the density $\exp(ivN_T + (1-e^{iv})A_T)$, see Giesecke (2007). The measure P^v neutralizes any feedback of events on the intensity λ. Formula (13) applies to all portfolio intensity specifications discussed in Section 3.2 above.

The expectation on the right hand side of equation (13) is a familiar expression in the defaultable term structure literature. It is analogous to the price at t of a security that pays $\exp(izY)$ at T if the issuer survives to T and 0 otherwise, assuming the issuer defaults at intensity $(1-e^{iv})\lambda$. The calculation of this price is well understood for a wide range of parametric intensity specifications, including affine and quadratic models. The reason is that this price is analogous to the price of a security paying $\exp(izY)$ at T in a default-free economy, where the short rate is $(1-e^{iv})\lambda$. Formula (13) thus extends the analytical tractability offered by existing term structure model specifications to the portfolio default process N.

To obtain the distribution of N we must invert the transform (13). To this end, for real a, b and x consider the conditional expectation

$$G_t(x; a, b, T, Y) = E\big[(a + bY)1_{\{N_T \leq x\}} \,|\, \mathcal{F}_t\big], \tag{14}$$

which is almost surely an increasing function in x that is constant on the intervals $[n, n+1)$ for n an integer, and vanishes for $x < 0$. The Fourier-Stieltjes transform of $G_t(x; a, b, T)$ can be obtained by integration by parts. For real v we get the formula

$$\mathcal{G}_t(v; a, b, T, Y) = \int_{-\infty}^{\infty} e^{ivx} dG_t(x; a, b, T, Y) = E\big[(a + bY)e^{ivN_T} \,|\, \mathcal{F}_t\big],$$

which can be expressed directly in terms of a partial derivative of the transform formula (13). For all non negative integers n we have

$$G_t(n; a, b, T, Y) = \frac{1}{2\pi} \int_{-\pi}^{\pi} \frac{e^{-ivn} - e^{iv}}{1 - e^{iv}} \, \mathcal{G}_t(v; a, b, T, Y)\, dv. \tag{15}$$

The inversion formula (15) recovers the conditional distribution function of N_T for $a = 1$ and $b = 0$. The function (14) can also be used to calculate the conditional default probabilities of the portfolio constituents. As explained in Section 3.2 above, in a top down model the constituent intensities are obtained by random thinning of the portfolio intensity λ. For any thinning process Z^k for firm k such that for $t \leq T$,

$$P\big[t < \tau^k \leq T \,|\, \mathcal{F}_t\big] = \int_t^T E\big[Z^k_s \lambda_s \,|\, \mathcal{F}_t\big] ds \tag{16}$$

see Giesecke & Goldberg (2005). The quantity $Z^k\lambda$ is the top down counterpart to the constituent intensity λ^k in a bottom up model, and formula (16) is the top down counterpart to the bottom up constituent probability formula (10).

Example 4.3. Let $\mathbb{F}$ be the filtration generated by the portfolio default process N. Suppose N is the Hawkes process considered in Errais et al. (2006), which is calibrated to the tranche market by Giesecke & Kim (2007). This model is a special case of Example 3.6 above. The portfolio intensity λ satisfies

$$\lambda_t = c + \delta \int_0^t e^{-\kappa(t-s)} dN_s. \tag{17}$$

The parameter $c > 0$ describes the base intensity. The parameter $\delta \geq 0$ governs the sensitivity of the intensity to defaults, and $\kappa \geq 0$ is the rate at which the impact of an event decays exponentially. Writing $d\lambda_t = \kappa(c - \lambda_t)dt + \delta dN_t$ shows that N is an affine point process, i.e. λ follows an affine jump diffusion process in the sense of Duffie, Pan & Singleton (2000) whose jump term is N. Further, a suitable version of the Girsanov-Meyer theorem implies that under the complex measure P^v, the point process N has intensity $e^{iv}\lambda$, see Giesecke (2007). Together with the transform formula (13), these observations allow us to conclude that

$$\begin{aligned} E\big[e^{iz\lambda_T + iv(N_T - N_t)} \,|\, \mathcal{F}_t\big] &= E^v\big[e^{iz\lambda_T - (1-e^{iv})\int_t^T \lambda_s ds} \,|\, \mathcal{F}_t\big] \\ &= e^{\alpha(t) + \beta(t)\lambda_t} \end{aligned} \qquad (18)$$

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

$$\partial_t \beta(t) = 1 + \kappa\beta(t) - e^{iv + \delta\beta(t)}$$

$$\partial_t \alpha(t) = -c\kappa\beta(t)$$

with boundary conditions $\beta(T) = iz$ and $\alpha(T) = 0$. By following the steps above, the transform (18) can be inverted to obtain the function $G_t(n; a, b, T, \lambda_T)$, which yields the distribution of the Hawkes process and is used to calculate the constituent default probabilities. To illustrate this, consider the thinning process of Example 3.8, given by

$$Z_t^k = \sum_{n=1}^{m} z^{kn} 1_{\{T^{n-1} < t \leq T^n\}}$$

where $(z^{kn})_{k,n=1,2,\ldots,m}$ is a matrix of non negative constants for which all rows and all columns sum to 1, and m is the number of portfolio constituents. In view of the default probability formula (16), it remains to calculate

$$E\big[Z_s^k \lambda_s \,|\, \mathcal{F}_t\big] = \sum_{n=1}^{m} z^{kn}\big(G_t(n-1; 0, 1, s, \lambda_s) - G_t(n-2; 0, 1, s, \lambda_s)\big).$$

The specification (17) does not guarantee that the portfolio intensity vanishes after the mth default in the portfolio. In other words, the process N governed by the intensity (17) can have more than m events. This is innocuous for typical portfolios, for example CDX index portfolios with more than 100 constituents. Here the distribution of the number of events is well approximated by the distribution of N. Nevertheless, it is straightforward to obtain the distribution of the stopped process $N^m = N \wedge m$ from that of N:

$$P\big[N_s^m - N_t^m = k \,|\, \mathcal{F}_t\big] = \begin{cases} P[N_s - N_t = k \,|\, \mathcal{F}_t] & \text{if} \quad k < m - N_t^m \\ P[N_s - N_t \geq k \,|\, \mathcal{F}_t] & \text{if} \quad k = m - N_t^m \\ 0 & \text{if} \quad k > m - N_t^m \end{cases}$$

Note that the truncation is not required for the constituent models, since the thinning process vanishes at the mth default by construction. □

5 Calibration

Accurate and stable intensity parameter estimation is a prerequisite for many applications of a portfolio credit model. For measurement and management of portfolio credit risk, the model is formulated under actual probabilities and must fit to historical default experience. For trading and hedging of standard or exotic portfolio derivatives, the model is formulated under risk-neutral probabilities and must fit to index and tranche market prices.

A specification of the joint evolution of constituent intensities can be fitted jointly to single- and multi-name market data. Eckner (2007), Feldhütter (2007) and Mortensen (2006) fit doubly stochastic models with jump diffusion risk factor dynamics to spreads of single name and tranche swaps observed on a given day. Papageorgiou & Sircar (2007) fit models with stochastic volatility risk factors. These papers obtain accurate fits.

There are two distinct ways to fit a top down model to market data. A specification of the portfolio intensity and the thinning processes can be calibrated jointly to single- and multi-name market data. In a procedure that does not require single name models or data, a stand-alone portfolio intensity specification can be fitted to index and tranche market data. Given the fit of the portfolio intensity, the constituent thinning processes can be calibrated to single name market data in an *optional* second step.

Most available empirical analyses fit a stand-alone specification of the portfolio intensity. Using different models, Arnsdorf & Halperin (2007), Brigo et al. (2006), Cont & Minca (2008), Ding et al. (2006), Giesecke & Kim (2007), Laurent et al. (2007) and Lopatin & Misirpashaev (2007) obtain accurate fits to index and tranche spreads of several maturities, all observed on a fixed date. With time-dependent parameters, the data can be matched perfectly. Azizpour & Giesecke (2008) and Longstaff & Rajan (2007) fit alternative portfolio intensity models with constant parameters to time series of index and tranche spreads for a fixed maturity. They find that the models replicate the time-series variation of market spreads for all tranche attachment points and maturities.

6 Conclusion

Dynamic reduced form models of portfolio credit risk provide many advantages over the static copula models that are in widespread use in the financial industry. First, they have realistic features that are motivated by empirical observation. Second, they specify the time evolution of the portfolio default process, and generate the portfolio loss distribution for all future horizons. Third, they accurately fit index and tranche market prices for all attachment points and maturities.

Dynamic reduced form models can be specified in two ways. In a bottom up model the constituent default intensities are the primitives. Such a specification is appropriate for the analysis of portfolios of highly heterogeneous constituents. It brings the information of the single-name market to bear on the calibration of the model. In a top down specification the portfolio default intensity is the modeling primitive, and constituent intensities are generated by random thinning. Since constituent calibration is optional for such a specification, a stand-alone portfolio intensity model can be used in situations with little or unreliable single name market information. An example is a reasonably granular portfolio of bonds or loans for which an index contract is traded. Even if single name information is available in principle, the sheer size of a portfolio can motivate the use of a stand-alone portfolio intensity model. Another application area for such a model is the analysis of exotic portfolio derivatives such as index and tranche forwards and options. These products are driven by the volatility of portfolio loss, which is conveniently controlled by the portfolio intensity.

References

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

Azizpour, Shahriar & Kay Giesecke (2008), Premia for correlated default risk. Working Paper, Stanford University.

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

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

Collin-Dufresne, Pierre, Robert Goldstein & Jean Helwege (2003), Are jumps in corporate bond yields priced? Modeling contagion via the updating of beliefs. Working Paper, Carnegie Mellon University.

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

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

Das, Sanjiv, Darrell Duffie, Nikunj Kapadia & Leandro Saita (2007), 'Common failings: How corporate defaults are correlated', *Journal of Finance* **62**, 93–117.

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

Delloye, Martin, Jean-David Fermanian & Mohammed Sbai (2006), Estimation of a reduced-form credit portfolio model and extensions to dynamic frailties. Working Paper, BNP Paribas.

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

Duffie, Darrell, Andreas Eckner, Guillaume Horel & Leandro Saita (2006), Frailty correlated default. Working Paper, Stanford University.

Duffie, Darrell & David Lando (2001), 'Term structures of credit spreads with incomplete accounting information', *Econometrica* **69**(3), 633–664.

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

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

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

Eckner, Andreas (2007), Computational techniques for basic affine models of portfolio credit risk. Working Paper, Stanford University.

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

Feldhütter, Peter (2007), An empirical investigation of an intensity based model for pricing CDO tranches. Working Paper, Copenhagen Business School.

Frey, Rüdiger & Jochen Backhaus (2004), Portfolio credit risk models with interacting default intensities: a markovian approach. Working Paper, Department of Mathematics, University of Leipzig.

Giesecke, Kay (2006), 'Default and information', *Jornal of Economic Dynamics and Control* **30**(11), 2281–2303.

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

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

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

Giesecke, Kay & Lisa Goldberg (2005), A top down approach to multi-name credit. Working Paper, Stanford University.

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

Giesecke, Kay & Stefan Weber (2004), 'Cyclical correlations, credit contagion, and portfolio losses', *Journal of Banking and Finance* **28**, 3009–3036.

Giesecke, Kay & Stefan Weber (2006), 'Credit contagion and aggregate loss', *Journal of Economic Dynamics and Control* **30**, 741–761.

Halperin, Igor (2007), Climbing down from the top: Single name hedging in top down credit models. Working Paper, Quantitative Research J.P. Morgan.

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

Jorion, Philippe & Gaiyan Zhang (2007*a*), Credit correlations from counterparty risk. Working Paper, University of California at Irvine.

Jorion, Philippe & Gaiyan Zhang (2007*b*), 'Good and bad credit contagion: Evidence from credit default swaps', *Journal of Financial Economics* **84**(3), 860–883.

Kim, Jack (2007), Time-inhomogeneous birth processes in credit modeling. Working Paper, Stanford University.

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

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

Laurent, Jean-Paul, Areski Cousin & Jean-David Fermanian (2007), Hedging default risks of CDOs in markovian contagion models. Working Paper, BNP Paribas.

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

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

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

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

Papageorgiou, Evan & Ronnie Sircar (2007), Multiscale intensity models and name grouping for valuation of multi-name credit derivatives. Working Paper, Princeton University.

Schönbucher, Philipp (2004), Information-driven default contagion. Working Paper, ETH Zürich.

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

Yu, Fan (2007), 'Correlated defaults in intensity based models', *Mathematical Finance* **17**, 155–173.