Simulating Correlated Defaults

Darrell Duffie and Kenneth Singleton*

May 21, 1999

Graduate School of Business, Stanford University

Abstract: This paper reviews some conventional and novel approaches to modeling and simulating correlated default times (and losses) on portfolios of loans, bonds, OTC derivatives, and other credit positions. We emphasize the impact of correlated jumps in credit quality on the performance of large portfolios of positions.

*Duffie is with the Graduate School of Business, Stanford University, Stanford, CA 94305, duffie@baht.stanford.edu. Singleton is with the Graduate School of Business, Stanford University, Stanford, CA 94305 and NBER, kenneths@future.stanford.edu. Our biggest acknowledgement is to Jun Pan for exceptional research assistance. We are grateful for conversations with Josh Danziger of CIBC and David Lando, and for motivating discussions with Gifford Fong of Gifford Fong Associates and Jim Vinci of PriceWaterhouseCoopers. This work is supported by PriceWaterhousCoopers Risk Institute, of which the authors are members, and by the Gifford Fong Associates Fund at the Graduate School of Business, Stanford University. The paper is downloadable at `http://www-leland.stanford.edu/~duffie/`.

Contents

1 Introduction

Computationally efficient methods for simulating default times for positions with numerous counterparties are central to the credit risk-management and derivative-pricing systems of major financial institutions. The likelihood of default of a given counterparty or borrower in a given time period is typically small. Computing the distribution of default times or losses on a large portfolio to reasonable accuracy may therefore require a significant number of simulated scenarios. This paper describes several computationally efficient frameworks for simulating default times for portfolios of loans and OTC derivatives, and compares some of the features of their implied distributions of default times.

Our focus is on the simulation of correlated credit-event times, which we can treat for concreteness as the default times of a given list of entities, such as corporations, private borrowers, or sovereign borrowers. To put the computational burden of a typical risk-management problem in perspective, consider a hypothetical portfolio consisting of 1,000 randomly selected firms rated Baa by Moody's, and suppose the risk manager is interested in 10-year scenarios. As indicated by the average default rates for 1970-97 in Figure 1, Baa firms experienced default at a rate of 0.12% per year, on average, over this period. Our sample portfolio of 1,000 Baa firms would thus have experienced a total average rate of approximately 12 defaults per year over this period. A "brute-force" simulation of default times for the portfolio using, say, weekly survival-default simulation would call for $10 \times 52 \times 1000 = 0.52$ million survive-or-default draws per 10-year scenario for this portfolio. One also must update, week by week, the conditional default probabilities of each entity, for another 0.52 million draws. Given random variation in exposures at default, we find that an order of magnitude of roughly 10,000 independent scenarios may be appropriate for estimation of "long-tail" confidence levels on total default losses for this sort of portfolio. (Reduction in the computational burden would likely come from variance-reduction or importance-sampling methods.) Simulation of 10,000 scenarios by the brute-force approach would thus call for on the 10 billion of survive-or-default random draws for each set of model parameters. In order to conduct stress testing calibration, many such exercises may be called for.

Fortunately such computationally intensive algorithms are unnecessary for many risk-management and pricing applications. Instead, one can use a variant of the following basic *recursive event-time* simulation algorithm for

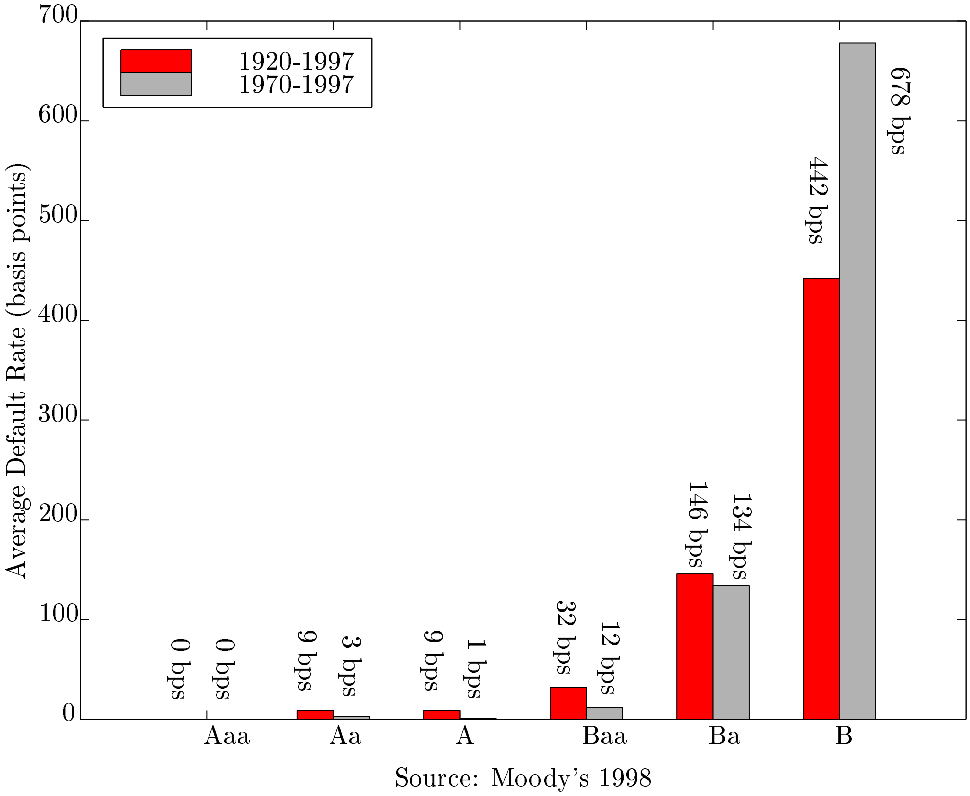

Figure 1: One year, weighted-average default rates by Moody's rating.

generating random multi-year scenarios for default times on a portfolio:

1. Given the simulated history to the last default time T_k, simulate the next time T_{k+1} of default of any entity. If T_{k+1} is after the lifetime of the portfolio, stop.

2. Otherwise, simulate the identities of any entities defaulting at T_{k+1}, as well as any other variables necessary to update the simulation model for the next default time.

3. Replace k with $k + 1$, and go back to Step 1.

Algorithms based on recursive event-time simulation are relatively efficient for large portfolios of moderate or low credit risk. For our hypothetical portfolio of 1,000 Baa counterparties, ignoring migration of credit quality for the moment, the recursive event-time algorithm would call for an average of about 12 random inter-default-time draws, and 12 draws for the identities of the defaulting entities, per 10-year scenario. We will present several frameworks that allow for random variation in an entity's credit-quality over time, while still allowing for the basic efficiency of the recursive event-time simulation algorithm. Moreover, recursive event-time simulation accommodates correlation among default times, including correlations caused by credit events that induce simultaneous jumps in the expected arrival rates of default of different counterparties.

For bank-wide risk management decisions, one may be interested in the likelihood that there will exist *some* interval of a given length, say 10 days, within the given multi-year planning horizon, during which default losses exceed a given amount of a bank's capital. This could be useful information, for example, in setting the bank's capital, structuring its portfolio for liquidity, or setting up provisional lines of credit. For accuracy in this calculation, it would be necessary to simulate the default times of the different entities to within relatively fine time slots, say daily.[1] Under the obvious proviso that the underlying probabilistic model of correlated default times is appropriate, we show that the recursive event-time algorithm is also well suited for this task, as it generates the precise default times implied by the model, scenario by scenario. When implemented for some hypothetical portfolios, we find that such measures as the distribution of losses for the "worst two weeks within 10 years" are particularly sensitive to one's assumption about correlation among entities.

2 Economic Framework

Throughout this analysis, we take as given a formulation of the default probabilities of each of the counterparties. We remain agnostic about the nature of the economic information driving default and the structural model linking this information to the default event. By way of background, we begin by presenting a brief discussion of the frameworks that one might consider when

[1] For example, given a loss on a certain day, a subsequent loss 11 days later should not enter the associated 10-day loss window, whereas a loss 9 days later should. If one is only interested in measures of default losses over fixed accounting periods, say quarterly, then the compuational burden of the brute-force simulation approach can be reduced somewhat by moving from daily to perhaps quarterly frequency for survival-default random draws.

modeling default events.

For corporate entities, foundational research on corporate debt pricing by Black and Scholes (1973) and Merton (1974) suggests modeling default probabilities at a given horizon as the likelihood that assets will not be sufficient to meet liabilities. Important variants and extenstions of the model allowing treatment of term-structures include those due to Geske (1977), Leland (1994), Leland and Toft (1996), and Longstaff and Schwartz (1995). This class of models is the theoretical underpinnings of the popular commercial "EDF" measure of default probability supplied by KMV Corporation.[2] That it has predictive power for rating migrations and defaults is shown, for example, by Delianedis and Geske (1998).

As far as implementation of this corporate-finance framework for purposes of simulating correlated default times for a list of firms, at least two methodologies are feasible:

A. Simulate the underlying asset processes until "first-passage" of assets to default boundaries, as follows:

   (a) From equity or bond (or both) price data, fit correlation and volatility parameters of log-asset-value processes, modeled, say, as Brownian motions.

   (b) Fit, as well, boundaries, possibly time dependent, that determine the default time of each firm as the first time that its assets cross its default boundary.

   (c) Simulate the sample paths of the underlying correlated Brownian motions, and record the first passage times for each firm.

B. The "intensity" of default for a given firm is the conditional expected arrival rate of default, which may depend on many observables. (A constant intensity implies a "Poisson arrival" of default.) One can fit stochastic default-intensity models to publicly traded borrowers, from market data on equity and bond prices, and perhaps other data such as transition histories or macroeconomic business-cycle variables.[3] Then

[2] See Kealhofer (1995).

[3] For example, Wilson (1997) proposes to fit transition intensities to country-level economic performance measures. This implies independence of default times for within-country entities, conditional on domestic business-cycle variables. One could model correlation induced by industry performance by a similar mechanism.

simulate correlated default times using well known algorithms, some reviewed below, for simulation of stopping times with given intensities.

A special case of the intensity-based Approach B is the CreditMetrics model of JP Morgan, for which one fits a model of credit-rating transitions, with correlations in rating changes that are implied by estimated correlations in changes in firm values. In effect, the default-time distribution for a given firm is based on historical rating transition and default data, while correlations in ratings changes can be "calibrated" to market cross-firm equity return correlations.[4]

In principle, Approaches A and B both allow one to use information about the prices of corporate securities when fitting a default-timing model. Use of this information seems desirable, because:

- Equity and corporate bond prices reflect market views of both asset valuations as well as asset volatilities. For individual entities, this provides enough information, in the context of the model, to derive "risk-neutral" default probabilities. Time-series data may supply risk premia with which one could estimate intensities.
- Correlation in the timing of default among different entities can be inferred through the model.

For the first-passage simulation method, Approach A, simulation of the paths of asset levels until first passage to default boundaries is burdensome for a large number of entities over long time horizons. Moreover, for long time horizons, it is not obvious how to control this simulation method for realistic re-capitalization possibilities. Calibration to reasonable long-term default probabilities may be intractable, if not unrealistic.

In the remainder of this paper, we focus on default-time simulation using Approach B, based on various models of default-arrival intensities. We

[4]The algorithm is roughly as follows: *(i)* Estimate the correlations of the various entities asset re-valuations, based on historical equity returns. *(ii)* For each entity, choose various levels of asset returns as cutoff boundaries for ratings or default so as to match the desired (say, historically estimated) rating transition probabilities. See, for example, *Introduction to Credit Metrics,* J.P. Morgan, New York, April 1997, page 26. By scaling arguments, volatilities are irrelevant given transition probabilities, unless one incorporates mean-return effects. *(iii)* Simulate a joint Gaussian vector for asset returns of all entities, and allocate the entities to their new ratings, or default, based on the outcomes.

emphasize, however, that to the extent permitted by the data, it is advantageous to fit the model of stochastic intensities of the entities to market equity or bond price behavior, economy-wide business cycle data, industry performance data, and the historical timing of ratings and default. Bijnen and Wijn (1994), Lennox (1998), Lundstedt (1999), McDonald and Van de Gucht (1996), and Shumway (1997) provide some examples. Intensity-based simulation is consistent in theory with first-passage simulation. For example, Duffie and Lando (1998) present a simple model in which, because of imperfect observation of a firm's assets (or the default boundary), there is a default intensity that can be estimated from asset reports.

3 Multi-Entity Default Intensity

This section reviews the basic ideas of intensity modeling for multi-entity portfolios. Appendix A reviews the more primitive underlying single-entity intensity model.

Throughout, we use the fact that the sum $N = N_r + N_R$ of independent Poisson arrival processes N_r and N_R, with respective intensities r and R, is itself a Poisson process with intensity $r + R$. The same property applies with randomly varying arrival intensities, under certain conditions, the most critical of which is that the arrivals cannot occur simultaneously. With simultaneous default, one can instead formulate separate credit events that could cause more than one entity to default at the same time, and to model the intensity of such joint credit events, as we shall see.

3.1 Constant and Independent Default Risk

Suppose, for illustration, that there are n_A counterparties of type A, each with default at intensity h_A, and n_B of type B, each with default intensity h_B. The time horizon T is fixed for this discussion. The intensities are assumed to be constant. We assume, for now, that two firms cannot default at the same time. In this case, we can simulate the defaulting firms, and times of their defaults, by the following simple version of the recursive event-time algorithm.

1. We first simulate a single Poisson process up to time T, with intensity $H = h_A n_A + h_B n_B$. For example, with $n_A = 1000$ A-rated counterparties, each with default intensities of $h_A = 0.001$ per year, and $n_B = 100$

B-rated counterparties, each with with an intensity of $h_B = 0.05$ per year, we have a total intensity of $H = 1 + 5 = 6$ defaults per year. The inter-arrival times T_1, and $T_i - T_{i-1}$, for $i > 0$, are simulated as independent exponentially distributed variables of parameter H, and successively added together to get the times $T_1, T_2, \ldots$. This gives us the default times and the number of firms defaulting before T.

2. At each default time T_i, the firm to default is of type A with probability $a = n_A h_A/(n_A h_A + n_B h_B)$, and is of type B with probability $1 - a$. We draw a random variable Y, called a "Bernoulli trial," whose only outcomes are A and B, with $P(Y = A) = 1 - P(Y = B) = a$. If the outcome of Y is A, we select one of the remaining type-A firms to default at time T_i, at random. (That is, each firm is drawn with probability $1/n_A$.) If the outcome of Y is B, we draw one of the remaining type-B firms, again at random.

3. We could adjust the arrival rates of default events of each type as firms default or otherwise disappear. This nuisance can be avoided simply by "not counting" firms that have already defaulted once. For some portfolios, such as revolving portfolios of bonds or loans underlying collateralized debt obligations, one can accomodate the introduction of new entities over time based on the simulated performance of the portfolio to date as well as simulated market data such as interest rates. One could extend the contagion model of Davis (1999) by modeling a jump in the default intensity of one entity with the default by some other entity.

In this way, we obtain, for each scenario, a list of the counterparties that default and times at which each defaults. Because, typically, only a small fraction of counterparties default in a given simulated scenario, this may be an efficient computational method for simulating total default losses. Relatively few market values and netting effects need to be computed.

3.2 Joint Credit Events

Certain credit events may be common to a number of counterparties. These could include:

- Severe catastrophes (for example, "Earthquake in Tokyo").

- Systemic default or liquidity breakdowns.
- Sovereign risks, such as a default, a moratorium on capital outflows, or a devaluation.
- Counterparties linked to each other through contracts or capital structure.

With joint credit events, some of the default intensity of each entity is tied to an event at which any entity may default, with some given probability. The total default intensity of entity i at time t is

$$h_{it} = p_{it} J_t + H_{it},$$

where, at time t,

- J_t is the intensity for arrival of joint credit events.
- p_{it} is the probability that entity i defaults given a joint credit event.
- H_{it} is the intensity of arrival of default specific to entity i.

With this model, the intensity of arrival of any kind of event is

$$H_t = J_t + H_{1t} + \cdots + H_{nt}.$$

In order to simulate defaults, one can adopt the following variant of the recursive event-time algorithm, which allows for correlation both through correlated changes in default intensities, as well as through joint credit events. For now, we will assume that default intensities are changing deterministically between defaults and credit events. Later, we allow them to be general correlated random processes, subject to technical conditions and of course the tractability necessary to apply our general algorithmic approach.

1. Generate the next credit event time t, conditional on current information, based on the total intensity process H. If $t > T$, stop.
2. At the event time t, allocate the event, as joint, with probability, $p_J(t) = J_t/(J_t + H_{1t} + \cdots + H_{nt})$, or not joint, with probability $1 - p_J(t)$.

3. If the type of event is simulated to be joint, then survive-or-default draws are simulated for each entity, independently or with correlation, depending on the model specification, entity i defaulting with probability p_{it}.

4. If the simulated event is not joint, then one of the counterparties is drawn at random to default. Entity i is drawn with probability $H_{it}/(H_{1t} + \cdots + H_{nt})$.

5. Any defaulting counterparties are deleted from the list.

6. According to the model specification, the intensity processes $H_1, \ldots, H_n, J$ and event-conditional default probabilities $p_1, \ldots, p_n$ are reset, conditioning on the history to date, and one returns to Step 1.

In the simplest version of the model, the only adjustment of intensities in the last step is to replace the intensity process H_i and event-conditional default probability p_i of any defaulting entity with zeros. More general versions are discussed below.

3.3 Example: Multivariate-Exponential Default Times

A special case of this approach is the classical multivariate-exponential distribution of failure times, reviewed in Appendix A, under which all individual intensities (H_i and J) and conditional default probabilities (p_i) are constant. In this case, each entity's default time is literally a Poisson arrival.

The main advantage of the multivariate exponential model is its simplicity. Simulation is easy. Numerous statistics (moments, joint survival probabilities, and so on) are easily calculated explicitly. For example, for first-to-default swap pricing, the impact of correlation is easily captured.

On the other hand, the model is not easily calibrated to data that bear on the term structure of default probabilities, such as bond or equity price data. For example, with risk-neutral bond investors, the term structure of credit yield spreads for the multivariate exponential model of default times is literally flat, because the default hazard rates are constant, whereas credit spreads often exhibit substantial slope, volatility, and correlation. (Term-structure effects could also be influenced by time-variation in conditional expected recoveries at default, as in Das and Tufano (1996), or in risk premia.)

Moreover, it is somewhat unrealistic to suppose that two or more firms would default literally simultaneously, unless there is a parent-subsidiary or similar contractual relationship. While the difference between simultaneous and nearly-timed default may not be critical for expected default losses or for the pricing of certain credit derivatives, it may indeed be an important distinction for measurement of the likelihood of a given sized loss within a given time window. With the multivariate exponential model, to the extent that correlations in the incidence of defaults within a given year are realistically captured, the model may imply an unrealistic amount of default within a given week or month.

4 Default Time Simulation Algorithms

Given stochastic intensity processes $h_1, \ldots, h_n$ for each entity, our objective is to simulate the associated default times $\tau_1, \ldots, \tau_n$, as well as the identities of the defaulter at each default time. Two of the basic algorithms, coming from the reliability literature on failure-time modeling,[5] are reviewed below. In some cases, as we shall point out, these algorithms are computationally burdensome for general correlated multi-variate diffusion models for intensities. In later sections, we suggest some tractable models.

4.1 Single-Entity Default Time Simulation

The primitive for the single-entity case is a stochastic intensity process h for a default time τ. As a first step, it is helpful if one can tractably compute, at any time t before default, the conditional probability $p(t, s)$ of survival to any given time $s > t$. Under certain conditions, the conditional survival probability $p(t, s)$ is given by

$$q_t(s) = E_t \left[\exp \left(\int_0^t -h(s) \, ds \right) \right], \tag{1}$$

where E_t denotes conditional expectation given the information available at time t. A key condition[6] for this result is that, for the fixed time horizon s, the process $\{q_s(t) : t \geq 0\}$ defined by (1) does not jump at τ. For example,

[5] See, for example, the survey by Shaked and Shanthikumar (1993).

[6] For more on this condition, see Duffie, Schroder, and Skiadas (1996). Kusuoka (1998) provides some interesting examples in which this condition fails.

it is enough that h is a diffusion process, or a jump diffusion with jump times that are not default times. We would want a model for h allowing $q_t(s)$ to be easily computed. Several such models are discussed in the following section.

In any case, there are two well known algorithms for simulation of the default time τ:

(A) **Inverse-CDF Simulation:** Build a model in which the survival probability $p(0,t)$ is easily calculated. Simulate a uniformly distributed random variable U, and let τ be chosen[7] so that $p(0,\tau) = U$.

(B) **Compensator Simulation:** Build a model in which the accumulated intensity, $H(t) = \int_0^t h(u)\,du$, often called the "compensator," is feasibly simulated. Simulate, independently of H, a standard (unit mean) exponentially distributed variable Z. Let τ be chosen[8] so that $H(\tau) = Z$.

Compensator simulation can be intractable unless the compensator can be easily simulated. For diffusion models of intensity, exact compensator simulation is relatively computationally intensive if the sample paths of the underlying diffusion must be simulated. One could use Euler or higher-order schemes for discrete-time approximate simulation of the stochastic differential equations underlying the intensities. Depending on the number of discrete time periods and the number of scenarios, this may be relatively expensive.

4.2 Multi-Entity Default Time Simulation

Suppose the event times $\tau_1, \ldots, \tau_n$ have respective intensity processes $h_1, \ldots, h_n$. For simplicity, simultaneous default is ruled out. (That is, the probability that $\tau_i = \tau_j$ is assumed to be zero for $i \neq j$. Otherwise, one reduces to simulation of the underlying event times at which simultaneous default may occur.)

One can simulate the times $\tau_1, \ldots, \tau_n$ with the correct joint distribution (including correlation of course) by either of the following basic algorithms, letting T denote the time horizon, assuming that one is interested in knowing the outcome of τ_i only if it is in the interval $(0, T)$.

[7]This assumes that $p(0,t) \to 0$ as $t \to \infty$. If not, then let $\tau = \inf\{t : p(0,t) \leq U\}$, which may have $+\infty$ as an outcome.

[8]This assumes that $H(t) \to \infty$ as $t \to \infty$. If not, then let $\tau = \inf\{t : H(t) \geq Z\}$, which may have $+\infty$ as an outcome.

(A) **Recursive Inverse-CDF Simulation:** Extending the single-entity algorithm, one proceeds as follows,[9] letting T_k denote the k-th to occur of the default times, I_k the identity of the k-th defaulter, $A(k)$ the set of undefaulted entities after T_{k-1}, and

$$W_k = \{(T_1, I_1, Y_1), \ldots, (T_k, I_k, Y_k)\},$$

the set of conditioning variables available at time T_k, where Y_k denotes a list of any additional state variables used for computing the CDF $p(T_k, \cdot)$ of the next default time T_{k+1}. For example, if intensities are assumed to functions of a Markov state variable X, then $X(T_k)$ would be included in Y_k. We let $h^{(k)} = \sum_{i \in A(k)} h_i$ denote the total intensity of default over the remaining entities, and we let

$$q_k(t, s) = E\left[\exp\left(\int_t^s -h^{(k)}(s)\, ds\right) \,\middle|\, W_{k-1}\right]. \tag{2}$$

We have $p(T_k, s) = q_k(T_k, s)$ for any $s > T_k$ provided that $q_k(\,\cdot\,, s)$ does not jump at T_{k+1}, which we assume. Under technical conditions, the conditional probability given (W_{k-1}, T_k) that i is the defaulting entity at T_k is

$$g_i(k) = P(I_k = i \,|W_{k-1}, T_k) = \frac{\gamma_i(T_{k-1}, T_k)}{\sum_{j \in A(k)} \gamma_j(T_{k-1}, T_k)},$$

where[10] for each $s > T_{k-1}$,

$$\gamma_i(T_{k-1}, s) = E\left[\exp\left(\int_{T(k-1)}^s -h^{(k)}(u)\, du\right) h_i(s) \,\middle|\, W_{k-1}\right]. \tag{3}$$

The steps of the algorithm are as follows.

1. Let $k = 1$, $T_0 = 0$, and $A_0 = \{1, \ldots, n\}$.
2. At time T_{k-1}, simulate, by inverse-CDF simulation using $p(T_{k-1}, \cdot)$, the next-to-default time T_k.
3. If $T_k > T$ stop.

[9] We leave out technical conditions and details.
[10] For more, but not all, details, see Duffie (1998a).

4. Simulate the identity I_k of the k-th defaulter from $A(k)$, with the conditional probability that $I_k = i$ equal to $g_i(k)$.
5. Simulate the additional "state" variables Y_k, with their distribution given the conditioning variables W_{k-1}, T_k, I_k.
6. Remove I_k from $A(k-1)$ to get $A(k)$, and unless $A(k)$ is empty, advance k by 1, and go back to Step 2.

(B) **Multi-Compensator Simulation:** Under technical conditions, the following algorithm generates stopping times[11] $\tau_1, \ldots, \tau_n$ with the given intensity processes $h_1, \ldots, h_n$. It is assumed that the compensator $H_i(t) = \int_0^t h_i(u)\, du$ can be simulated for each i and t.

(a) Simulate n independent unit-mean exponentially distributed random variables $Z_1, \ldots, Z_n$.

(b) For each i, if $H_i(T) < Z_i$, then $\tau_i > T$.

(c) Otherwise, let $\tau_i = \min\{t : H_i(t) = Z_i\}$.

The compensator simulation approach is also possible, in a more complicated form, if one computes the intensities when conditioning only the history of the default times and identities of defaulting entities. This information structure is sometimes called the "internal filtration," and the resulting intensities in this setting are often called conditional hazard rates. The failure-time simulation is then called the "multivariate hazard construction," proposed[12] by Norros (1986) and Shaked and Shanthikumar (1987). The multivariate-hazard construction is preferred if the hazard rates relative to the internal filtration can be computed explictly.

For our numerical results in later sections, we use recursive inverse-CDF simulation.

5 Conventional Stochastic Intensity Models

A wide range of models of stochastic intensity processes with a significant level of analytical tractability have been used for modeling individual default

[11] There is no claim of uniqueness. Uniqueness in distribution is however implied if one assumes that, conditional on the paths of the intensities $h_1, \ldots, h_n$, the stopping times $\tau_1, \ldots, \tau_n$ are independent.

[12] This is based on a result of Meyer (1971).

times. We will briefly consider some of these. We later propose some alternative models that are relatively simple and tractable for correlated default-time simulation. As we shall see, they have rather different implications for the impact of correlation.

We note the helpful analogy between survival probabilities and discount functions, as it is apparent from (1) that, for risk-neutral bond investors, $q_t(s)$ is mathematically equivalent, under technical conditions, to a zero-coupon bond price at time t, for maturity s, at short interest rate h. This suggests convenient classes of intensity processes that have already been used to model interest rates.

5.1 Diffusion Intensities

A convenient diffusion model of intensities would take $h_i(t) = f_i(X_t)$, for some state-space process X of a simple, say Cox-Ingersoll-Ross (1985) (CIR) or more generally affine,[13] form. Because the "discount" $q_t(s)$ is explicit for this family of models, provided f_i is itself affine, it is simple to simulate individual entity default times. One can likewise simulate the first-to-default time $T_1 = \min(\tau_1, \ldots, \tau_n)$ by the inverse-CDF method as the CDF of T_1 is easily computed.[14] In order to simulate subsequent chronologically ordered default times $T_2, T_3, \ldots$, there are different approaches. For example, suppose T_1 and the first entity I_1 to default have been simulated as in the recursive inverse-CDF algorithm above. In order to simulate T_2, the second default time, one could first simulate the state vector $X(T_1)$ determining intensities at time T_1, and then given $X(T_1)$, compute the new CDF $p(T_1, \cdot)$ for the second default time T_2, and so on. The conditional distribution of $X(T_1)$ given (T_1, I_1) can be computed easily for CIR models[15] although it need not be easily simulated in multivariate CIR cases. Compensator simulation is feasible through simulation of the sample paths of X, but may be computationally intensive. One could alternatively undertake compensator simulation using the multivariate hazard construction, that is, with respect to the internal filtration, although the detailed hazard-rate calculations have not yet been worked out for non-trivial cases beyond the CIR model, to our knowledge.

[13] See Duffie and Kan (1996).
[14] See, for example, Duffie (1998a).
[15] See Duffie (1998a).

5.2 Finite-State Continuous-time Markov Chains

Jarrow, Lando, and Turnbull (1996) proposed the use of finite-state continuous-time Markov chains as a model of intensities, for purposes of term-structure modeling for defaultable bonds. They identified states with credit ratings. Arvanitis, Gregory, and Laurent (1998), Lando (1998b), Li (1998), Nakazato (1998), and Tepla (1999) have extended the model to allow stochastic variation of intensities within rating.

As for multi-entity correlated rating-transition models, Lando (1998a) considered correlation within the framework of finite-state continuous-time Markov chains for each entity's intensity, by taking each of the states to correspond to a particular list of ratings, by entity. For example, with 2 entities and 3 ratings, A, B, and C, the states are

$$\{AA, AB, AC, BA, BB, BC, CA, CB, CC\},$$

where, for example, AB is the state for which entity 1 is currently in rating A and entity 2 is in rating B. This leads to an exponential growth in the size of the state space as the number of entities and ratings grow.

For a more tractable version of this approach, one could assume symmetry conditions among entities, and that transition times for all entities within one rating to any other particular rating or default are multivariate exponential. Given the first time of a default or transition out of a particular rating, symmetry calls for drawing the identity of the defaulting or transitioning entities by picking at random, equally likely, some entity currently in that rating. One then re-sets the total-intensity model for various transitions, draws another transitition time, and so on. A typical element of the state space is simply a list consisting of the number of entities in each rating.[16]

5.3 HJM Forward Default Probabilities

One can formulate Heath-Jarrow-Morton (1992) (HJM) style models of the term structure of survival probabilities, by again exploiting the analogy between bond prices and survival probabilities. One formulates the conditional probability at time t of survival to time s as the process

$$p(t, s) = \exp\left(-\int_t^s f(t, u)\, du\right), \tag{4}$$

[16] The state space is $\{1, \ldots, n\}^K$, where n is the number of entities and K is the number of ratings.

where for each fixed s, one supposes that $f(\cdot\,,s)$ is an Ito process. One can add the jumps to the formulation of f. From the model specified for f, and the HJM "drift" restriction[17] imposed on f by the fact that $\{p(t,s):\, t \geq 0\}$ is a martingale, one obtains indirectly a stochastic model for the intensity process $h(t) = f(t,t)$. The default time τ can then by simulated, for example by inverse-CDF or compensator simulation.

HJM-style models of survival probabilities could be used to simulate correlated default times for the various entities by compensator simulation, although this may be compuationally burdensome.

6 Correlated Jump Intensity Processes

One example that we have explored is a model in which individual entities have default intensities that mean revert, with correlated Poisson arrivals of randomly sized jumps. By formulating the individual-entity default intensity jump times as multivariate exponential, one arrives at a relatively simple but useful model for simulating correlated defaults.

6.1 Mean-Reverting Intensities with Jumps

First we formulate a single entity's intensity process as a mean-reverting process with jumps. Specifically, the intensity process h of a typical entity's default time has independently distributed jumps that arrive at some constant intensity λ, and otherwise h mean reverts at rate k to a constant θ. A sample path for such an entity, suffering 4 modest jumps to its intensity, is illustrated in Figure 2. For this illustration, and our examples to follow, the mean-reversion rate is $k = 0.5$. One can easily generalize.

With this simple model, the default arrival intensity $h(t)$ of the default

[17]Suppose, for each fixed s, that $df(t,s) = \mu(t,s)\,dt + \sigma(t,s)\,dB_t$, where B is a standard Brownian motion in $\mathbb{R}^d$ and μ and σ satisfy certain measurability and integrability conditions. Then, under additional technical conditions, the martingale property of $p(\cdot\,,s)$ implies that $\mu(t,s) = \sigma(t,s)\cdot\int_t^s \sigma(t,u)\,du$, for $t \leq s$. This is often called the "HJM drift restriction." See Heath, Jarrow, and Morton (1992) for the drift restriction on f, given its volatilities, and Duffie (1998b) for more details on the application to default risk. It is literally the case that $f(t,s)$ is the hazard rate for default at s, conditional on information at t. Simulation of f therefore makes for easy updates of conditional survival probabilies.

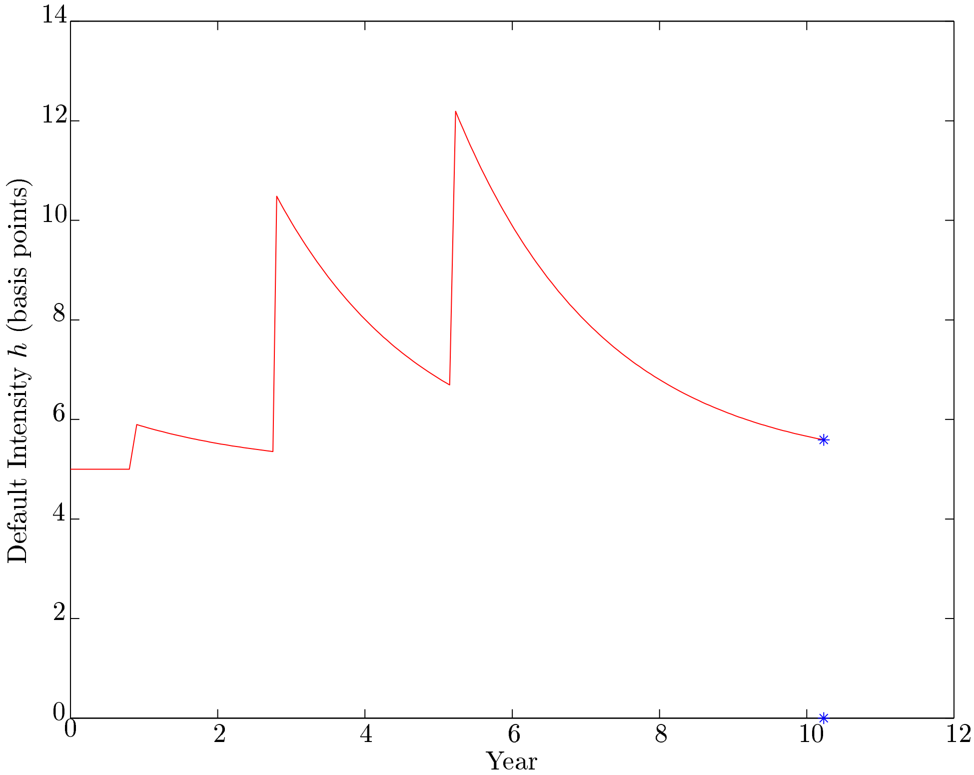

Figure 2: A simulated path of default intensity.

time τ, in between jump events, satisfies the ordinary differential equation

$$\frac{dh(t)}{dt} = k(\theta - h(t)). \tag{5}$$

Thus, at any time t between jump times, we have a simple solution

$$h_t = \theta + e^{-k(t-T)}(h_T - \theta), \tag{6}$$

where T is the time of the last jump and h_T is the post-jump intensity at time T.

For example, suppose that jumps in intensity are exponentially distributed with mean J. The initial condition $h(0)$ and the parameters (k, λ, J, θ) determine the probability distribution of the default time. In fact, it can be shown[18] that the conditional probability at any $t < \tau$ of survival from t to s is

$$p(t,s) = e^{\alpha(s-t)+\beta(s-t)h(t)}, \tag{7}$$

where

$$\begin{aligned} \beta(t) &= -\tfrac{1-e^{-kt}}{k} \\ \alpha(t) &= -\theta\left(t - \tfrac{1-e^{-kt}}{k}\right) - \tfrac{\lambda}{J+k}\left(Jt - \ln\left(1 + \tfrac{1-e^{-kt}}{k}J\right)\right). \end{aligned}$$

For example, suppose $\lambda = 0.001$, $k = 0.5$, $\theta = 0.001$, $J = 5$, and $h(0) = 0.001$, meaning an initial mean arrival rate of default of once per thousand years (10 basis points). For comparison, the average rate of default arrival for both A-rated and Aa-rated corporate issuers from 1920 to 1997 was 9 basis points, according to Moody's,[19] as illustrated in Figure 1.

At these parameters, a jump in default risk is likely to be devastating, as a mean jump in intensity of 5 implies a mean expected remaining life of less that 3 months. This model is slightly less risky[20] than one in which an issuer defaults at a constant intensity of 20 basis points. (For reference, the average default arrival rate for all Baa-rated corporate issuers for 1920 to 1997, as measured by Moody's, is 32 basis points, as indicated in Figure 1.)

So-called "risk-neutral" versions of this calculation can also be used as part of a term-structure model for defaultable debt, for example to calibrate parameters or to price credit derivatives. Suppose for simplicity that there are no credit risk premia. (The parameters (k, θ, λ, J) could be adjusted to account for risk premia, in a "risk-neutral" version of the model.) The t-year credit yield spread $S(t)$ for zero-recovery instruments (say bond coupons) is then given by

$$S(t) = -\frac{\alpha(t) + \beta(t)h(0)}{t}. \tag{8}$$

---

[18]The relevant ordinary differential equations for α and β are easily found, and then solved, given the conjectured form of p.

[19]See "Historical Default Rates of Corporate Bond Issuers, 1920-1997," Moody's Investor Services, Global Credit Research, February, 1998.

[20]This comparison follows from the fact that the jump-intensity model, at these parameters, starts an entity with a total arrival rate of 20 basis for a default *or* a potentially-survivable jump in intensity.

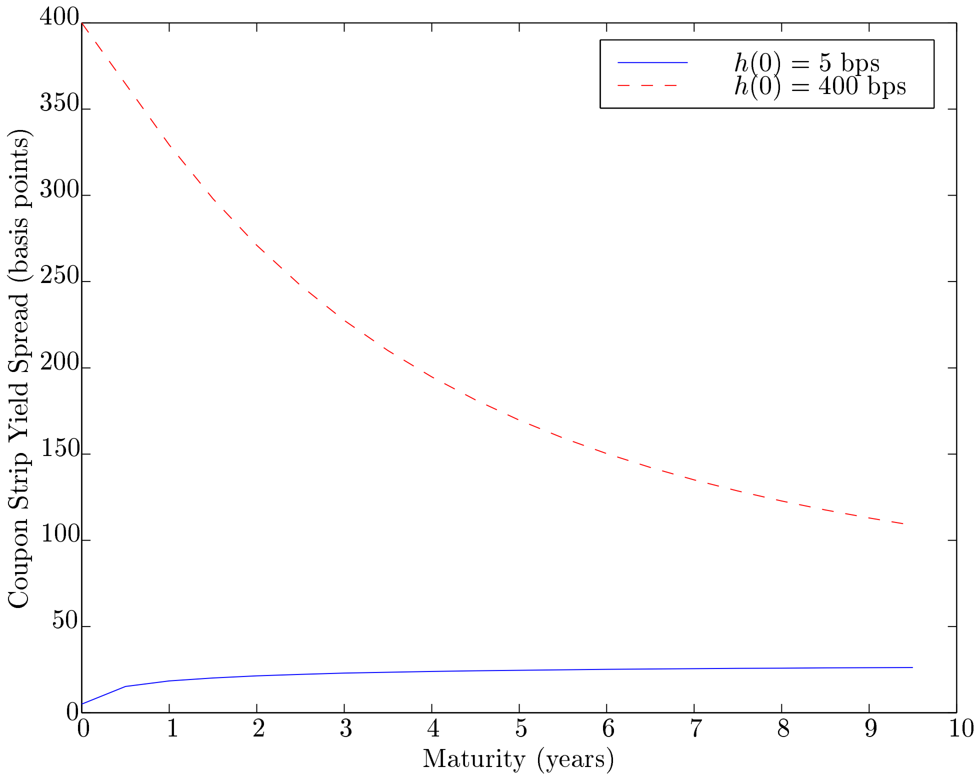

Figure 3: Term structure of coupon-strip (zero-recovery) yield spreads.

Credit spreads for a low-risk and a high-risk issuer are plotted in Figure 3. The two issuers have the same parameters ($\theta = 10$ basis points, $k = 0.5$, $J = 5$, and $\lambda = 10$ basis points). The low-risk issuer, however, has an initial default intensity $h(0)$ of 5 basis points. The high-risk issuer has an initial default arrival rate of 400 basis points. (For reference, the average default rates of B-rated corporate issuers for 1920-97, as measured by Moody's, was 442 basis points.)

In summary, this jump-intensity model is appealing on grounds of simplicity and tractability. As we shall see, it is also tractable and appealing as a foundation for modeling correlation in default times among various entities. In order to capture the effects of daily volatility in yield spreads (and quality), one can easily extend to multiple jump types, at different respective arrival rates, or affine dependence of h on an "affine" diffusion state variable. All of the above calculations can be extended to this case, giving easily calculated survival probabilities and default-time densities.[21]

6.2 Correlated Jump Intensity Processes

Suppose that entities $1, \ldots, n$ have correlated multivariate exponential event times, not for default, but rather for the times of sudden jumps of default-arrival intensities, in the context of the jump-intensity model just described. Once the individual parameters $(k_i, \theta_i, \lambda_i, J_i)$ are fixed, the only parameters to be chosen are those determining correlation across the multivariate-exponential jump times of the individual entity's intensities.

This model is particularly tractable for simulation of successive defaults, as all intensities (for individual entities' default times, for the arrival times of jumps in intensities, for the arrival of any default, and so on) are deterministic between credit event times (jumps in intensities or defaults). Thus, the next event time and the identity of the next event, conditional on the simulation to date, both have explicit cumulative distribution functions (CDFs), and can therefore be simulated by two indepedendent uniform-$[0, 1]$ draws. The first draw determines the time τ of the next credit event using the explicit inverse-CDF. The second draw determines the identity of the event, conditional on τ, as a jump time or a default time, and if a default, which entity defaulted. These various events have conditional probabilities proportional to their respective intensities at τ, which are in turn explicitly determined

[21] See Duffie, Pan, and Singleton (1998) for details.

from (6).[22]

This general class of intensity models is a special case of multivariate "affine" jump-diffusion intensity models, for which survival probabilities as well as moments, Laplace transforms, and Fourier transforms of default times and state variables can be computed by analytical means.[23]

6.3 Numerical Example

For illustration, consider an example in which entities' default times have the same parameters (k, θ, λ, J) determining the distributions of their individual default times. Suppose, for simplicity, that the sizes of jumps to intensities are independent (across time and counterparties) and exponentially distributed. It would not be difficult to allow for multivariate exponential jump sizes among entities affected by simultaneous jumps in intensities.

For a simple symmetric version of this model, we suppose that all correlation in the intensity jump times arises from a common Poisson process N_c, with intensity Λ_c. There are also "idiosyncratic" Poisson processes $N_1, \ldots, N_n$, with common intensity parameter Λ. At the k-th jump time T_k of N_c, for any k, the default intensity h_i of entity i jumps by an exponentially distributed amount Y_{ik} with mean J, if and only if $U_{ik} = 1$, where U_{ik} has outcomes 1 and 0 with probabilities p and $1-p$, respectively. At the k-th jump time T_{ik} of N_i, there is (with probability one) an exponentially distributed jump Z_{ik} in h_i with mean J. All of these event times and jump sizes are independent.[24]

The parameters Λ_c, p, and Λ are chosen to provide a given amount of correlation (within the limits imposed by the model structure), maintaining

[22] If the mean-reversion rates $k_1, \ldots, k_n$ are identical and equal to k, then the total current intensity $h(t) = \sum_{i \in A(t)} h_i(t)$, where $A(t)$ is the set of surviving firms at time t, satisfies piecewise, between credit event and default times, the same ODE (5), taking $\theta = \sum_{i \in A(t)} \theta_i$. Even without the common mean-reversion assumption, it can be seen with a few calculations that, at any time, given the current intensities, the time-to-next default or jump event has an explicit probability distribution and can therefore be simulated by the inverse-CDF method.

[23] See Duffie, Pan, and Singleton (1998).

[24] To be precise, $W = \{N_c, N_1, \ldots, N_n, \{U_{ik}, Y_{ik}, Z_{ik} : 1 \leq i \leq n,\ k \geq 1\}\}$ are independent, and conditional on W, the default times $\tau_1, \ldots, \tau_n$ have the conditionally deterministic intensity processes $h_1, \ldots, h_n$ just described. In fact, conditional on N_c, the default times are independent, and one can therefore simulate default times by first simulating the jump times of N_c, and then generate default times independently given N_c.

the arrival intensity

$$\lambda = p\Lambda_c + \Lambda \tag{9}$$

of a given entity's jumps in intensity.

Referring to Appendix A, we calculate the correlation between the multivariate-exponential jump times of any 2 entities to be

$$\rho = \frac{p^2\Lambda_c}{\Lambda_c(2p - p^2) + 2\Lambda}. \tag{10}$$

It will be noted that as J goes to infinity, the model approaches the multivariate exponential default-time model. As p approaches 0, the model converges to one of independent default intensities. As p converges to 1 and Λ converges to 0, the model approaches one of perfectly correlated jump intensities.[25]

For our base case, we take the parameters for individual-entity default risk to be those,

$$(\theta = 0.001, \lambda = 0.002, k = 0.5, J = 5),$$

used in our previous individual-entity illustration. For correlation parameters, we take

$$(p = 0.02, \Lambda_c = 0.05), \tag{11}$$

so that the rate of arrival of "idiosyncratic" jumps in an entity's default intensity is

$$\Lambda = \lambda - p\Lambda_c = 0.001.$$

This implies, for example, that the probability that h_i jumps at t given that h_j jumps at t is $p\Lambda_c/(\Lambda_c + \Lambda) = 1\%$.

For this base-case model, we simulated 20,000 independent scenarios for the correlated default times of 1000 entities over a 10-year period. This means that, in effect, we simulated defaults covering 200 million entity-years.[26]

[25] For perfectly correlated (that is, identical) default times, take $p = 1$, $\Lambda = 0$, and let $J \to \infty$, so that all entities default at the first jump in N_c, and not otherwise.

[26] The total CPU time expended for a Sun UltraSparc processor was approximately 3 hours. The software was written in MatLab. Simulation was pseudo-random, with no variance-reduction techniques.

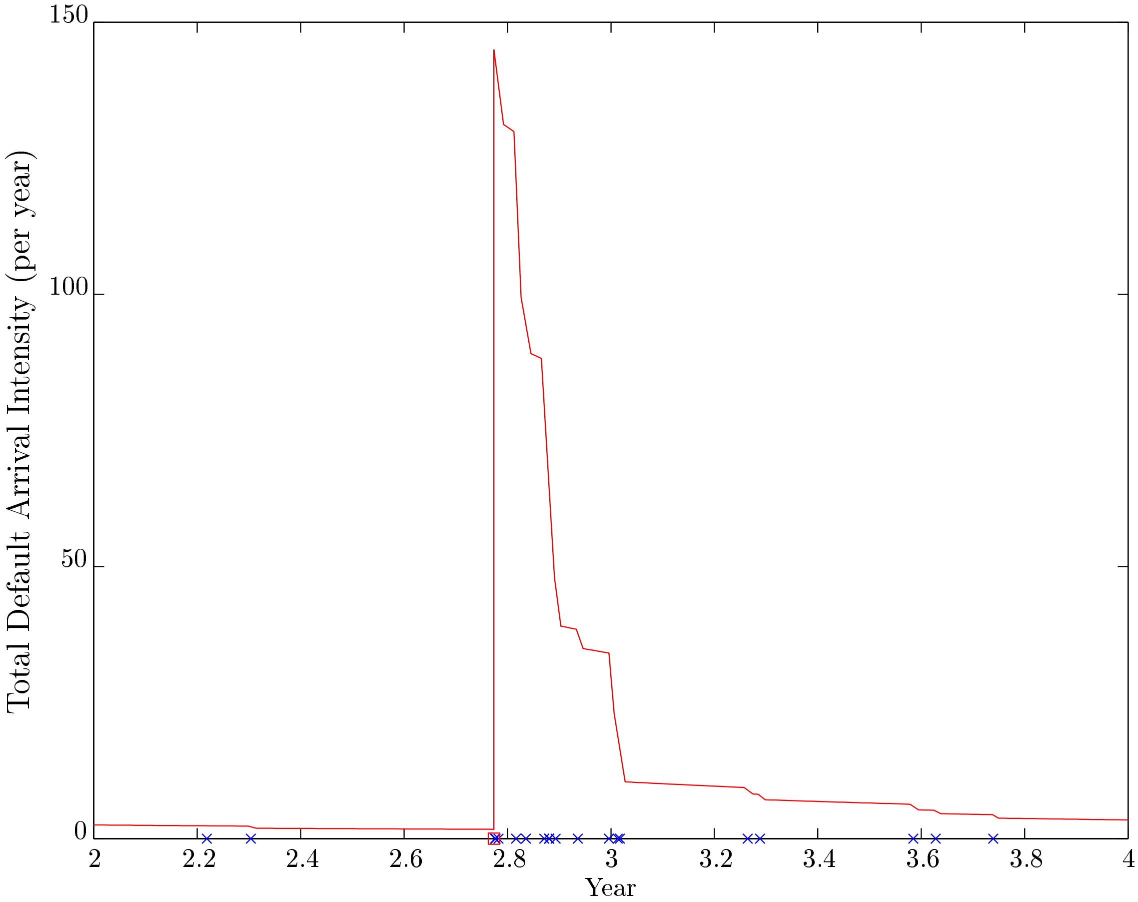

Figure 4: A portion of a simulated sample path of total default arrival intensity (initially 1,000 firms). A square indicates a market-wide credit event. An `x` indicates a default event.

A portion of a typical sample path for the total arrival intensity h of defaults for the 1000 original entities for this base-case model is illustrated in Figure 4. Along the horizontal (calendar-time) axis, a "box" is marked to show the arrival time of a jump in N_c, which on this sample path instigated (at random) jumps in default intensity for a number of entities. As some of these entities default, at times indicated by the symbol "$\times$" on the horizontal axis, and the intensities of default for the surviving firms revert back to typical levels, the total arrival intensity h for defaults drops quickly, moving back near its pre-event levels within roughly one year.

Opinions one may have about the reasonableness of the illustrated behavior may suggest adjustment of the parameters. Of course, ideally, the parameters could be fit to price, rating, default, and other data. For example, after adding risk premia to Λ_c, Λ, and J, one could calibrate to credit spreads using the explicit spread formula (8), after adjusting the parameters for risk premia.

Fixing individual default-time distributions, we consider variations in the correlation structure:

- Zero correlation ($\Lambda_c = 0$).
- "High" correlation ($\Lambda_c = 0.1, p = 0.02$). This implies that $\Lambda = 0$, and therefore that the probability that h_i jumps at t given that h_j jumps at t is 0.02. Higher correlation is of course possible by reducing Λ_c and increasing p, holding λ constant.

The probability of experiencing at least n defaults of the original 1000 firms in a particular quarter is shown in Figure 5, where we pick for illustration the first quarter of the 5th year. Other quarters show similar results, as indicated in an appendix.

Perhaps more telling, from the point of view of the impact of correlation on credit-risk management and measurement, is the likelihood of the existence of some m-day interval during the entire T-year time horizon during which at least n entities defaulted. For a time horizon of $T = 10$ years, $n = 4$ entities, and time intervals of various numbers of days (m), the results for the uncorrelated, base case, and high-correlation models are shown in Figure 6. Additional results are found in Appendix B.

These figures reveal the impact of correlation, albeit in the limited context of this model, for the ability of a given pool of bank capital to support a

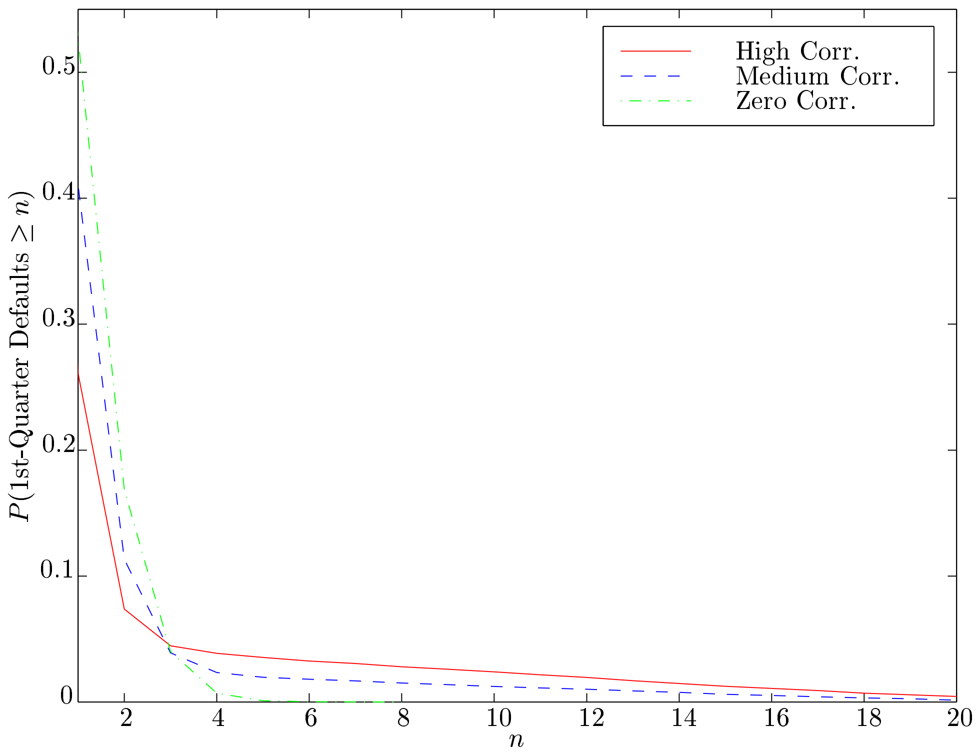

Figure 5: Probability of n or more defaults in the first quarter of year 5.

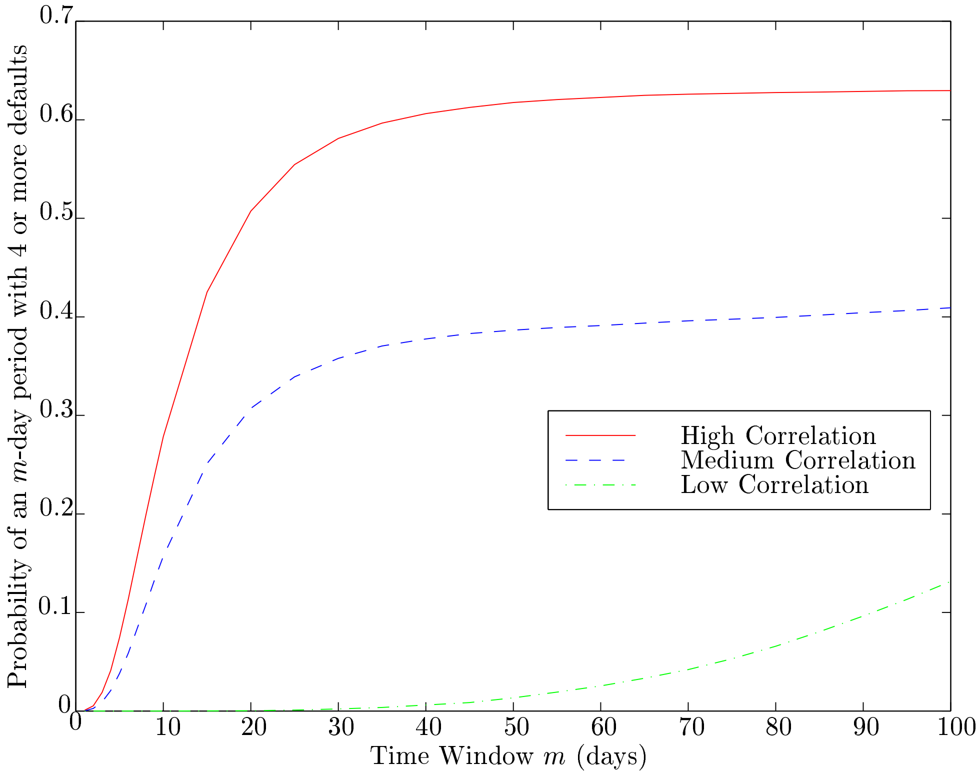

Figure 6: Probability of an m-day interval within 10 years having 4 or more defaults (base case).

given level of credit risk when it is anticipated that firm-specific or market-wide illiquidity shocks may prevent re-capitalization of a bank within a time window of a given size.

7 Correlated Log-Normal Intensities

In this section, we illustrate another simple example of multivariate correlated intensities. In this example, intensities are log-normally distributed and, for computational tractability, are assumed to be piece-wise constant, changing at a given frequency, say once per quarter or once per year.

For example, given the sparse data available and statistical fitting advantages, it may not be unreasonable to assume that, for counterparty i, the intensity h_{it} in year i is generated by a log-normal model with mean reversion, given by

$$\log h_{i,t+1} = \kappa_i(\log \overline{h}_i - \log h_{i,t}) + \sigma_i \epsilon_{i,t+1},$$

where

- κ_i is a rate of mean reversion.
- $\log \overline{h}_i$ is the steady-state mean level for the log-intensity.
- σ_i is the volatility of intensities.
- $\epsilon_{i,1}, \ldots, \epsilon_{i,t}$ is an independent sequence of standard-normal random variables.

One can introduce correlation in default risk through the correlation ρ_{ij} between ϵ_{it} and ϵ_{jt}.

7.1 Numerical Example

For this piece-wise log-normal model, we simulated default times and losses upon default, assuming that exposures are independent and exponentially distributed.[27]

[27]Exposures could equally well be simulated as the positive parts of joint-log-normals, correlated with the underlying intensities, at little or no additional computational burden. One could also take multivariate exponential exposures, also allowing for correlation.

We take two classes of firms. Type A firms, of which there are n_A, each have a mean exposure of 100 million dollars. Type B firms, of which there is a smaller number $n_B = n - n_A$, each have a mean exposure of 10 million dollars. All exposures are independent of each other and of the intensities.

The base case for our simulation is specified as follows. We take intensity update intervals of $D = 1$ year, $n = 5000$ firms, $n_A = 4500$, a time horizon of $T = 10$ years, mean reversion parameter $\kappa_i = 0.5$, intensity volatility parameter $\sigma_i = 0.4$. For type-A entities, we take $h_i(0) = \overline{h}_i = 0.0005$. For type-$B$, we take $h_i(0) = \overline{h}_i = 0.005$. We take the base-case intensity-shock correlation parameter $\rho_{ij} = 0.50$ for all i and j.

Figures 7 and 8 show the estimated median, 75%, 95%, and 99% confidence levels on default losses and number of defaults for the first quarter of each year, for years 1 through 10. The results are shown for 10,000 independent scenarios, requiring a total workstation cpu time of approximately 100 minutes.[28]

7.2 Check on Sampling Error

In order to check the Monte Carlo sampling errors of the reported confidence levels for default losses, we reduce the number of scenarios to 5,000, and conducted 10 independent samples of this type. Tables 1 and 2, found in Appendix C, report the sample means and standard deviations of selected confidence intervals for default loss and number of default events. At 5,000 scenarios, the estimated standard deviation of the sampling error of even the 99% confidence levels for total default losses in a fixed quarter are under 4 percent of those respective confidence intervals. The estimated sampling errors of the median losses are slightly higher in fractional terms, but much lower in absolute terms.

7.3 Check on Time-Discretization "Error"

While the model is not necessarily to be treated as a discrete-time approximation of a continuous-time intensity model, it could be. In that case, it makes sense to shorten the discretization interval D and re-estimate the loss distribution, taking the underlying log-intensity model to be Ornstein-Uhlenbeck

[28] The software was written in Fortran 90. The CPU time is for a single Sun UltraSparc processor. We used pseudo-random independent sampling with no variance-reduction methods.

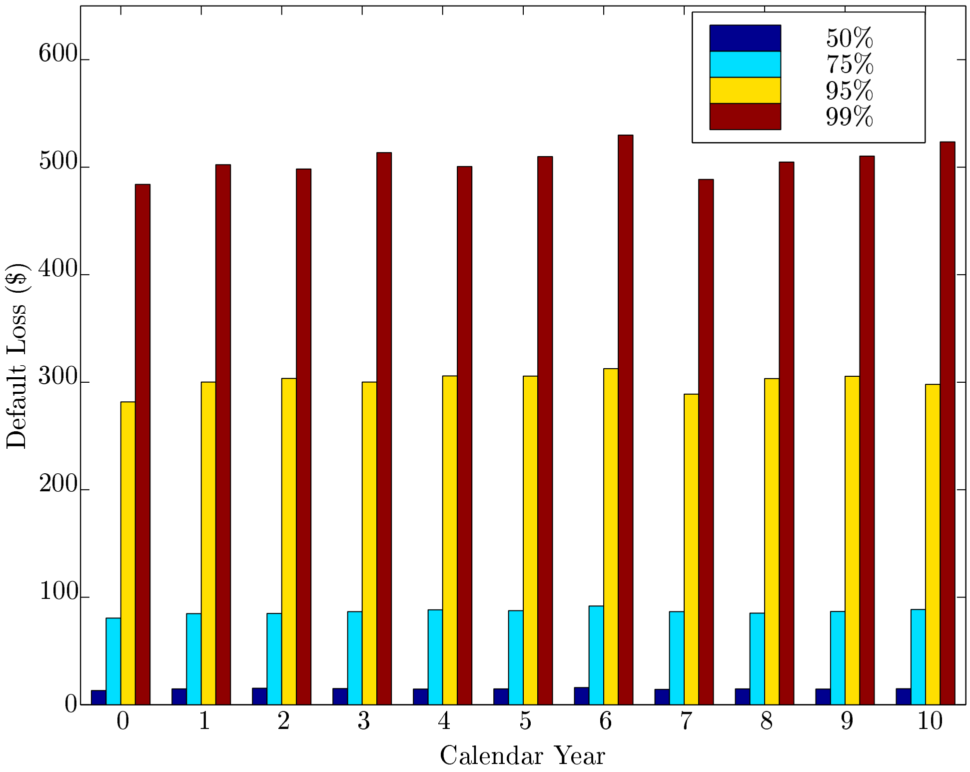

Figure 7: Simulated 50%, 75%, 95%, and 99% condidence levels on default losses for the first quarter of each year. 10,000 simulation runs for the base case.

driven by Brownian motion, controlling for fixed annual mean-reversion and variance parameters.

We shorten D from 1 to 0.5 and 0.25 years, keeping all else as in the base case. Again, 10,000 simulation runs are used. The total cpu times are 140 minutes and 240 minutes, respectively. The reader can review the results in Table 3, found in Appendix C. The "discretization error" seems reasonably small in light of parameter uncertainty, for this particular model. (One notes that, because of the Ornstein-Uhlenbeck model underlying the discretization, the log intensities at the beginning of each year have a fixed multivarite Gaussian distribution that is unaffected by discretization in this setting, an advantage of this model.)

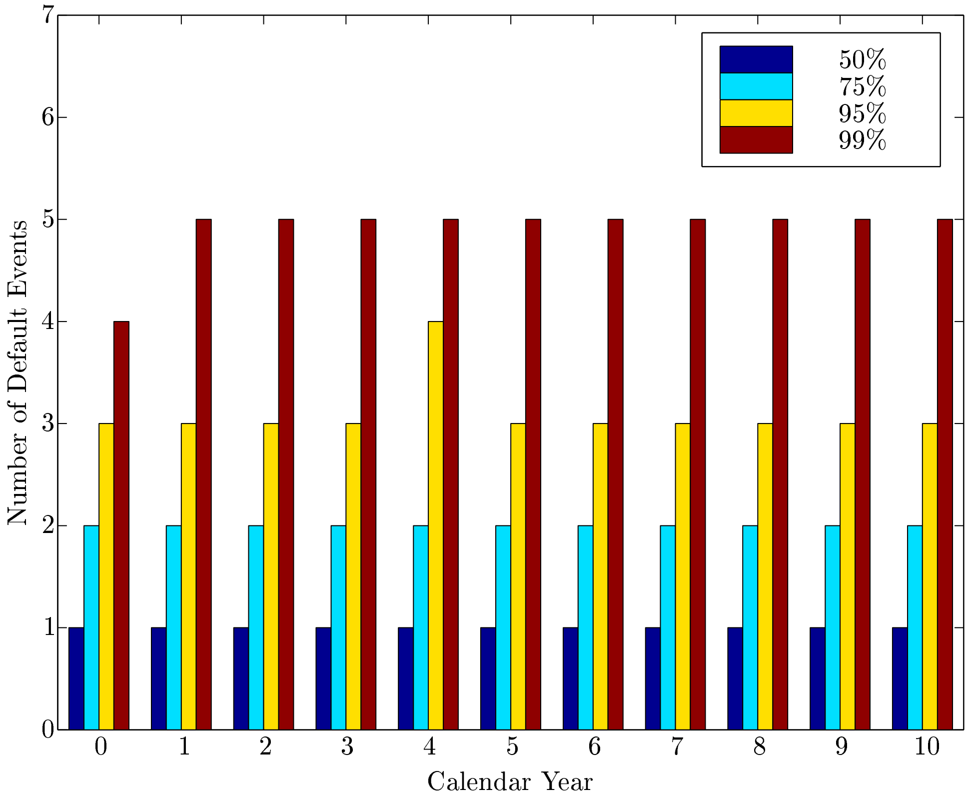

Figure 8: Simulated 50%, 75%, 95%, and 99% condidence levels on number of default events for the first quarter of each year. 10,000 simulation runs for the base case.

7.4 Role of Intensity Volatility and Correlation

The impacts of variation in the intensity volatility paramter σ_i (holding all else equal) confidence levels of default losses is shown in Figure 11, found in Appendix C.

The response of even high-confidence-level default losses to variation of the correlation parameter ρ_{ij} is relatively small at our base-case volatility, in comparison with the illustrated impact of correlation in the jump-intensity model. This insensitivity is illustrated in Figure 9. In order to obtain significantly higher impact of correlation, we apply a 100% volatility parameter, σ_i, with unconditional default probabilities roughly on a par with those at base case in our jump-intensity model. The results appear in Figures 12 and

13, found in Appendix C.

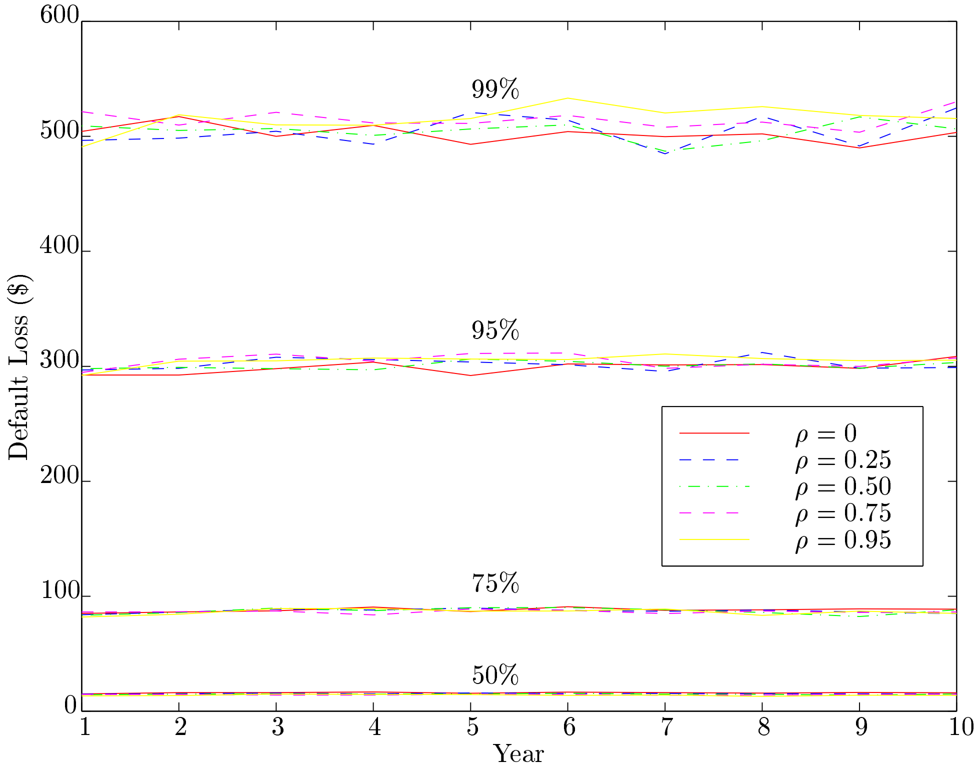

Figure 9: Default Loss for the first quarter of each year. The four bands correspond to 50%, 75%, 95%, and 99%-percentile default losses. Within each band, the n-percentile default losses associated with five different correlations of default intensity are marked by different colors. 20,000 simulation runs for the base case.

References

[1] A. Aravantis, J. Gregory, and J.-P. Laurent (1999) "Building Models for Credit Spreads," *Fixed Income Journal.*

[2] R. Barlow and F. Proschan (1981) *Statistical Theory of Reliability and Life Testing,* Silver Spring, Maryland.

[3] E. J. Bijnen and M.F.C.M. Wijn (1994), "Corporate Prediction Models, Ratios or Regression Analysis?" Working Paper, Faculty of Social Sciences, Netherlands.

[4] F. Black and M. Scholes (1973), "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* **81:** 637-654.

[5] P. Brémaud (1980) *Point Processes and Queues – Martingale Dynamics*, New York: Springer-Verlag.

[6] J. Cox, J. Ingersoll, and S. Ross (1985) "A Theory of the Term Structure of Interest Rates," *Econometrica* **53:** 385-408.

[7] S. Das and P. Tufano (1997) "Pricing Credit-Sensitive Debt When Interest Rates, Credit Ratings, and Credit Spreads are Stochastic," *Journal of Financial Engineering* **5:** 161-198.

[8] M. Davis (1999) "Contagion Modeling of Collateralized Bond Obligations," Working Paper, Bank of Tokyo Mitsubishi.

[9] G. Delianedis and R. Geske (1998) "Credit Risk and Risk Neutral Default Probabilities: Information About Rating Migrations and Defaults," Working Paper 19-98, Anderson Graduate School of Business, University of California, Los Angeles.

[10] D. Duffie (1998a) "First-to-Default Valuation," Working Paper, Graduate School of Business, Stanford University.

[11] D. Duffie (1998b) "Defaultable Term Structure Models with Fractional Recovery of Par," Working Paper, Graduate School of Business, Stanford University.

[12] D. Duffie and R. Kan (1996) "A Yield-Factor Model of Interest Rates," *Mathematical Finance* **6:** 379-406.

[13] D. Duffie and D. Lando (1997) "Term Structures of Credit Spreads with Incomplete Accounting Information," Working Paper, Graduate School of Business, Stanford University.

[14] D. Duffie, J. Pan, and K. Singleton (1998), "Transform Analysis and Option Pricing for Affine Jump Diffusions," Working Paper, Graduate School of Business, Stanford University.

[15] R. Geske (1977) "The Valuation of Corporate Liabilities as Compound Options," *Journal of Financial and Quantitative Analysis* **:** 541-552.

[16] D. Heath, R. Jarrow, and A. Morton (1992), "Bond Pricing and the Term Structure of Interest Rates", *Econometrica* **60:** 77-106.

[17] R. Jarrow, D. Lando, and S. Turnbull (1997), "A Markov Model for the Term Structure of Credit Spreads," *Review of Financial Studies* **10:** 481-523

[18] S. Kealhofer (1995) "Managing Default Risk in Portfolios of Derivatives," in E. Van Hertsen and P. Fields, editors, *Derivative Credit Risk, Advances in Measurement and Management*, London: Risk Publications.

[19] S. Kusuoka (1998), "A Remark on Default Risk Models," Working Paper, Graduate School of Mathematical Sciences, University of Tokyo.

[20] D. Lando (1997), "Modeling Bonds and Derivatives with Default Risk," *in* M. Dempster and S. Pliska (editors) *Mathematics of Derivative Securities*, pp. 369-393, Cambridge University Press.

[21] D. Lando (1998a), "On Rating Transition Analysis and Correlation," Working Paper, Department of Operations Research, University of Copenhagen, Forthcoming in *Credit Derivatives*, edited by R. Jameson, London: Risk Magazine.

[22] D. Lando (1998), "On Cox Processes and Credit Risky Securities," *Review of Derivatives Research.*

[23] H. Leland (1994), "Corporate Debt Value, Bond Covenants, and Optimal Capital Structure," *Journal of Finance* **49:** 1213-1252.

[24] H. Leland and K. Toft (1996), "Optimal Capital Structure, Endogenous Bankruptcy, and the Term Structure of Credit Spreads," *Journal of Finance* **51:** 987-1019.

[25] C. Lennox (1997) "Indentifying Failing Companies: A Re-evaluation of the Logit, Probit and DA Approaches," Working Paper, Economics

Department, Bristol University.

[26] T. Li (1997) "A Model of Pricing Defaultable Bonds and Credit Ratings," Working Paper, Olin School of Business.

[27] F. A. Longstaff and E. S. Schwartz (1995), "A Simple Approach to Valuing Risky Debt," *Journal of Finance* **50:** 789-821.

[28] K. G. Lundstedt and S. A. Hillgeist (1998), "Testing Corporate Default Theories," Working Paper, UC Berkeley and Northwestern University

[29] R. Merton (1974), "On The Pricing of Corporate Debt: The Risk Structure of Interest Rates," *The Journal of Finance* **29:** 449-470.

[30] P.-A. Meyer (1971) "Demonstration Simplifiée d'un Théorem de Knight," *Seminaire de Probabilités V*, Berlin: Springer.

[31] C.G. McDonald and L.M. Van de Gucht (1996), "The Default Risk of High-Yield Bonds," Working Paper, Louisiana State University.

[32] D. Nakazato (1997) "Gaussian Term Structure Models with Credit Rating Classes," Working Paper, Industrial Bank of Japan.

[33] I. Norros (1986) "A Compensator Representation of Multivariate Life Length Distributions, with Applications," *Scandinavian Journal of Statistics* **13:** 99-112.

[34] M. Shaked and J.G. Shanthikumar (1987) "The Multivariate Hazard Construction," *Stochastic Processes and Their Applications* **24:** 241-258.

[35] M. Shaked and J.G. Shanthikumar (1993) "Multivariate Conditional Hazard Rate and Mean Residual Life Functions and Their Applications," Chapter 7 of *Reliability and Decision Making*, edited by R. Barlow, C. Clarotti, and F. Spizzichino, London: Chapman and Hall.

[36] T. Shumway (1996) "Forecasting Bankruptcy More Efficiently: A Simple Hazard Model," Working Paper, University of Michigan Business School

[37] L. Tepla (1999) "On The Estimation of Risk-Neutral Transition Matrices for Pricing Credit Derivatives," Working Paper, INSEAD.

[38] T. Wilson (1997) "Portfolio Credit Risk," *RISK* **10** (October), pp. 56-61.

Appendices

A Review of Intensity Modeling

This appendix reviews the basic idea of default arrival intensities.

A.1 Poisson Arrival

A useful basic model of default for a given counterparty is one of Poisson arrival at a constant arrival rate, called "intensity," often denoted λ. For a given Poisson with *intensity* h,

- the probability of default over the next time period of small length Δ is approximately Δh.
- the probability of survival without default for t years is e^{-ht}.
- the expected time to default is $1/h$.

For example, at a constant intensity of 0.04, the probability of default within one year is approximately 4 percent, and the expected time to default is about 25 years.

The intensity of arrival of an event, in this sense is sometimes called the "hazard rate," which is more formally defined as $f(t) = -p'(t)/p(t)$, where $p(t)$ is the probability of survival to t, assuming that p is differentiable. The hazard rates are sometimes called "forward probabilities" in finance, and may be thought of as the intensities for a setting in which the only information resolved over time is the arrival of default. That is, $f(t)$ is the arrival rate of default at time t, conditional on no other information other than survival to t. Indeed, with constant intensity, the two terms, "hazard rate" and "intensity," are synonomous, as the time to default is exponentially distributed with parameter equal to intensity. This terminology varies.[29]

[29] To be more precise yet, the intensity λ of the point process N which starts at zero and jumps to one at the time itself, staying there indefinitely, is defined by the property that $\{N(t) - \int_0^t \lambda(s)\, ds : \quad t \geq 0\}$ is a martingale, after fixing the probability measure and filtration of σ-algebras defining information. See Brémaud (1980) for technical details. Thus, the intensity must drop to zero at the arrival time. We will speak loosely of the intensity of a Possion arrival to be a "constant" λ, even though the intensity drops to zero after arrival. This loose terminology makes sense if one speaks of intensity at t for t before arrival.

The classic Poisson model is based on the notion of independence of arrival risk over time. For example, the Poisson arrival at intensity h is approximated, with time periods of small length Δ, by the first time that a coin toss results in "Heads," given independent tosses of coins, one each period, with each toss having a probability $h\Delta$ of Heads and $1 - h\Delta$ of Tails. This "coin-toss" analogy highlights the unpredictable nature of default in this model. Though we may be an instant of time away from learning that an issuer has defaulted, when default does occur, it is a surprise.

A.2 Intensity Process

In practice, of course, as time passes, one would want to update the intensity for default by a given counterparty with new information that bears on the credit quality of that counterparty, beyond simply survival. That is, though the default event cannot be fully anticipated, the probability that one assigns to default will likely change over time unexpectedly. How much this probability changes over time depends on the available information about the issuer's financial condition and the reason for the default.

A natural model is to treat the arrival intensity, given all current information, as a random process. Assuming that intensites are updated with new information at the beginning of each year, and are constant during the year, it can be shown that the probability of survival for t years is

$$E[e^{-(h_0+h_1+h_2+\cdots+h_{t-1})}]. \tag{A.1}$$

In other words, looking forward from today (date 0), (A.1) gives the probability that the issuer will survive for t years. It is the probability of surviving the first year, times the probability of surviving the second year given the first year was survived, times the probability of surviving the third year given that the issuer survived until the second year, and so on. For a quarterly-update model, taking an annualized intensity of h_t at time t, the probability of survival for t years is

$$E\left[e^{-\frac{1}{4}(h_0+h_{0.25}+h_{0.5}+\cdots+h_{t-0.25})}\right]. \tag{A.2}$$

For a continuous-time model, under certain conditions,[30] we have the

[30]The conditions are reviewed in the text. Interesting exceptions are provided by Kusuoka (1998).

survival probability

$$S(t) = E\left[\exp\left(-\int_0^t h(s)\,ds\right)\right].$$

One can see an analogy between an intensity process h and a short interest rate process r: survival probability is to intensity as discount (zero-coupon bond price) is to short rate. In this analogy, the parallel to the "in-t-for-1 forward interest rate" is the "in-t-for 1 forward default rate," which is the probability of default in the period from t to $t+1$, conditional on no default before t. For example, in the constant-intensity model, the in-t-for 1 forward default rate is the intensity itself. Analogously, with a constant short interest rate, the forward rates are all equal to the short rate itself.

A.3 Multivariate Exponential Event Times

The simplest of all models of correlated credit event times is the multivariate exponential. The basic idea of the model is that all types of events, whether joint or particular to a given entity, have constant intensities. That is, with this model, each credit event still has a Poisson arrival with constant intensity, but certain entities may be affected simultaneously, with specified probabilities.

There are equivalent ways to specify such a model. The following is unconventional but convenient for applications to credit pricing and risk measurement.

The basic ingredients are independent Poisson processes $N_1, \ldots, N_m$ with intensity parameters $\lambda_1, \ldots, \lambda_m$. Whenever, for any j, there is a jump in process N_j, entity i has a credit event provided the outcome of an independent 0-or-1 trial, with probability p_{ij} of 1, is in fact 1.

We can think of the jumps of some (or all) of the underlying Poissons $N_1, \ldots, N_m$ as market-wide events that could, at random, affect any of n entities. Correlation effects are determined by the underlying credit-event arrival rates $\lambda_1, \ldots, \lambda_m$ and by the "impact" probabilities p_{ij}.

With this model, the arrivals of credit events for a given entity i is Poisson with intensity

$$g_i = \sum_{j=1}^{m} p_{ij}\lambda_j.$$

The intensity of arrival of simultaneous credit events for entities i and k is

$$g_{ik} = \sum_{j=1}^{m} p_{ij} p_{kj} \lambda_j.$$

Likewise, for any subset $A \subset \{1, \ldots, n\}$ of entities, the Poisson arrival rate of a simultaneous credit event for all entities in A is

$$g_A = \sum_{j=1}^{m} \lambda_j \prod_{i \in A} p_{ij},$$

where $\prod_{i \in A} p_{ij}$ denotes the product of p_{ij} over all i in A.

Conditional on survival of entities i and j to the current date, the correlation between the times to the next credit events for entities i and k turns out to be[31]

$$\rho_{ik} = \frac{g_{ik}}{g_i + g_k - g_{ik}}.$$

Many other statistics regarding the joint distribution of event times can be worked out explicity.[32]

Of course, for a credit event such as default, it makes sense to delete any defaulting entities from the model at default times. After such a time, the model for the timing of credit events of the remaining entities remains multivariate exponential.[33]

B More on the Jump Intensity Model

This appendix provides additional results regarding the jump-intensity model.

[31] See Barlow and Proschan (1981), p.135, Exercise 8(c). We are grateful to Josh Danziger of CIBC for bringing this convenient formula to our attention.

[32] See Barlow and Proschan (1981).

[33] An alternative for updating the model is to assume that the underlying Poisson processes "disappear" at certain (say Poisson) arrival times, and perhaps that others, with different parameters, "appear" at certain times. In this case, it is easy to update the model parameters with each appearance and disappearance, so that the model is piecewise-in-time multivariate exponential. Simulation in this framework is easily accomplished. First one simulates the appearance and disappearance times, which form "epochs." Then one simulates the event times within each epoch as exponentially distributed, with right censoring.

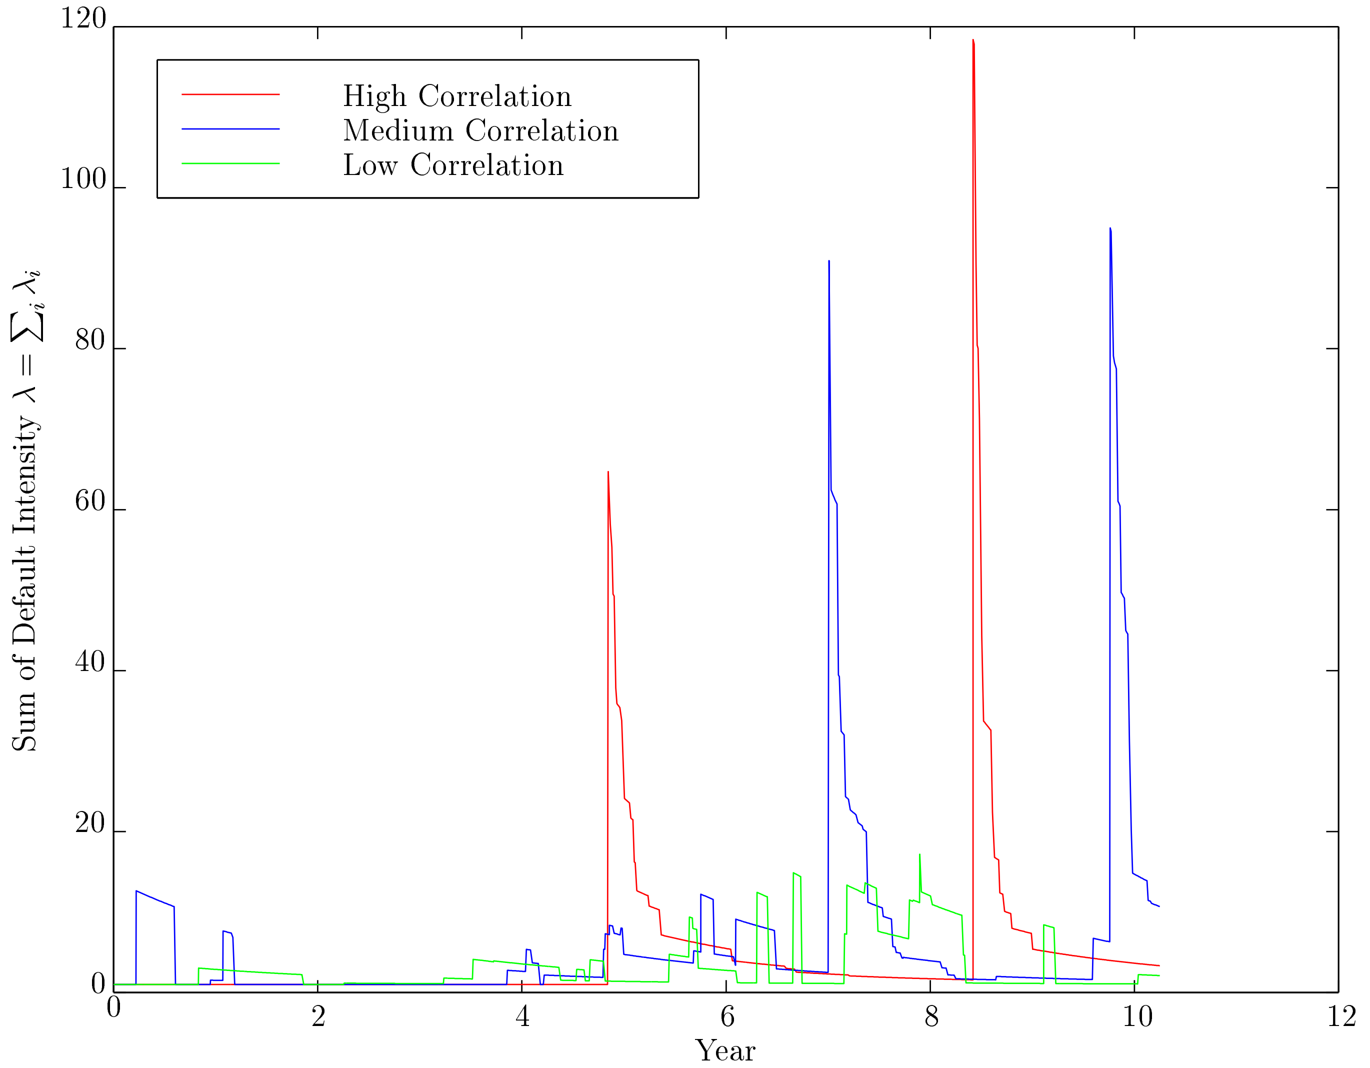

Figure 10: Sample paths of total default intensity, summed over 1,000 firms, for different correlations.

C More on the Log-Normal Intensity Model

Additional results for the log-normal intensity model are provided in this appendix.

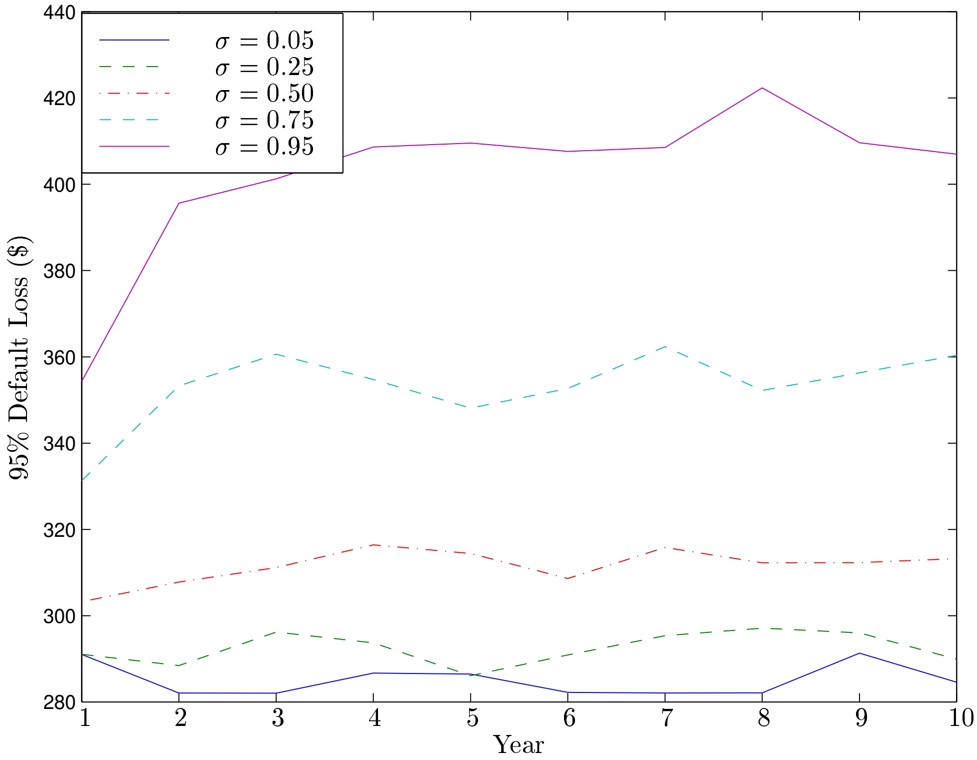

Figure 11: 95% Default Loss for the first quarter of each year. 20,000 simulation runs for the base case with varying volatility of default intensity.

Table 1: Sample mean and standard deviation (in parenthesis) of Simulated Decile Estimates for the 1st Quarter Default Losses, in Millions of Dollars. 5,000 simulations for each set of estimates. Sample statistics calculated from 10 sets of estimates.

Year	Confidence Levels			
	50%	75%	95%	99%
0	13.26	77.45	284.99	485.50
	(0.61)	(2.60)	(6.53)	(10.83)
1	14.76	85.00	302.04	513.06
	(0.57)	(3.18)	(12.48)	(21.19)
2	15.05	87.65	300.96	511.89
	(0.98)	(4.96)	(7.66)	(17.65)
3	15.02	86.25	299.90	511.19
	(0.54)	(2.42)	(8.96)	(21.25)
4	15.41	89.03	303.72	512.80
	(0.52)	(2.97)	(13.47)	(18.66)
5	15.11	87.45	303.50	513.02
	(0.89)	(4.48)	(7.30)	(13.90)
6	15.33	88.37	306.74	521.69
	(0.70)	(2.75)	(5.37)	(21.84)
7	15.04	86.77	299.95	509.42
	(0.73)	(4.75)	(8.32)	(16.04)
8	14.57	86.03	301.01	509.09
	(0.63)	(3.08)	(8.18)	(14.79)
9	14.91	90.05	306.48	512.63
	(1.04)	(2.50)	(8.17)	(22.52)
10	14.92	88.01	303.01	514.17
	(0.48)	(3.25)	(8.70)	(19.24)

Table 2: Sample mean and standard deviation (in parenthesis) of Simulated Decile Estimates for the 1st Quarter Default Invents. 5,000 simulations for each set of estimates. Sample statistics calculated from 10 sets of estimates.

Year	Confidence Levels			
	50%	75%	95%	99%
0	1.00	2.00	3.00	4.00
	(0.00)	(0.00)	(0.00)	(0.00)
1	1.00	2.00	3.00	4.80
	(0.00)	(0.00)	(0.00)	(0.42)
2	1.00	2.00	3.10	5.00
	(0.00)	(0.00)	(0.32)	(0.00)
3	1.00	2.00	3.40	5.00
	(0.00)	(0.00)	(0.52)	(0.00)
4	1.00	2.00	3.50	5.00
	(0.00)	(0.00)	(0.53)	(0.00)
5	1.00	2.00	3.10	5.00
	(0.00)	(0.00)	(0.32)	(0.00)
6	1.00	2.00	3.00	5.00
	(0.00)	(0.00)	(0.00)	(0.00)
7	1.00	2.00	3.00	5.00
	(0.00)	(0.00)	(0.00)	(0.00)
8	1.00	2.00	3.10	5.00
	(0.00)	(0.00)	(0.32)	(0.00)
9	1.00	2.00	3.00	5.00
	(0.00)	(0.00)	(0.00)	(0.00)
10	1.00	2.00	3.00	5.00
	(0.00)	(0.00)	(0.00)	(0.00)

Table 3: Simulated Decile Estimates for the 1st Quarter Default Losses, in Millions of Dollars. 10,000 simulations runs.

Year	Confidence Levels			
	50%	75%	95%	99%
$D = 1.0$ year				
0	13.26	80.56	281.64	483.94
2	15.39	84.88	303.60	498.33
4	14.61	88.26	305.90	500.67
6	15.95	91.86	312.63	529.85
8	14.81	85.19	303.39	504.73
10	14.89	88.65	298.07	523.53
$D = 0.5$ year				
0	14.01	80.99	288.12	485.95
2	15.42	89.15	304.10	499.60
4	14.91	87.72	294.60	512.09
6	14.88	88.76	304.69	483.56
8	13.65	84.31	311.54	526.44
10	14.32	85.58	302.32	521.78
$D = 0.25$ year				
0	13.18	79.61	288.65	517.70
2	15.38	90.34	304.36	517.58
4	14.66	85.50	309.05	522.20
6	15.22	87.00	300.71	504.99
8	15.43	91.28	314.25	529.74
10	15.26	90.37	302.40	513.56

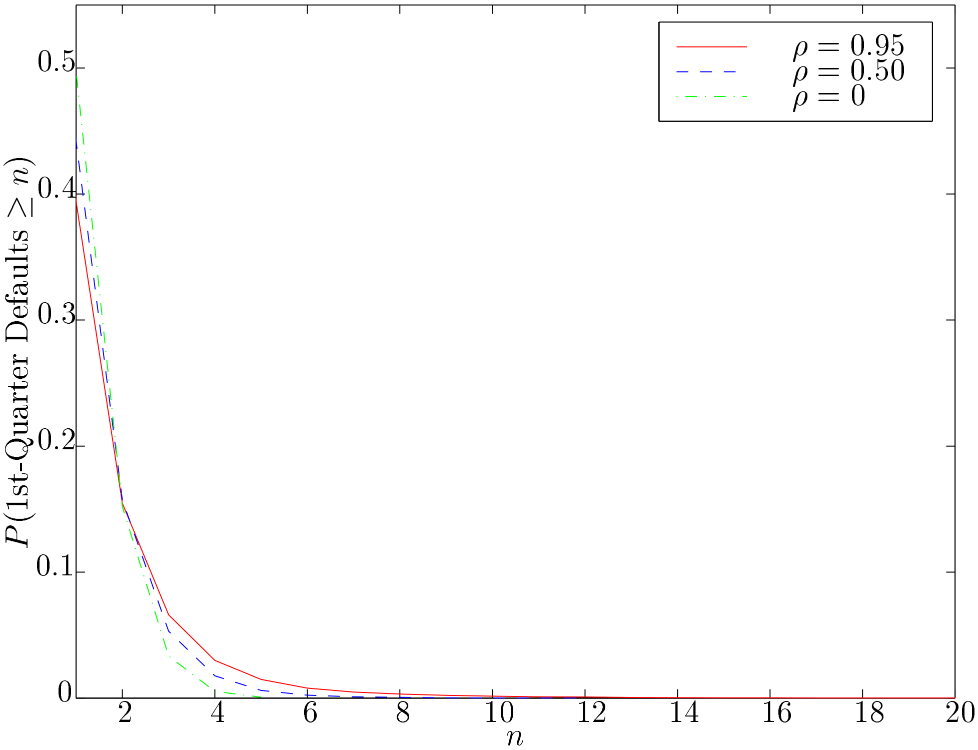

Figure 12: Probability of n or more defaults in the first-quarter of year 10. (1,000 entities, intensity exponential Ornstein-Uhlenbeck, parameters $\theta = \ln(0.0017)$, $\sigma = 1$, $\kappa = 0.5$, pair-wise intensity shock correlation ρ).

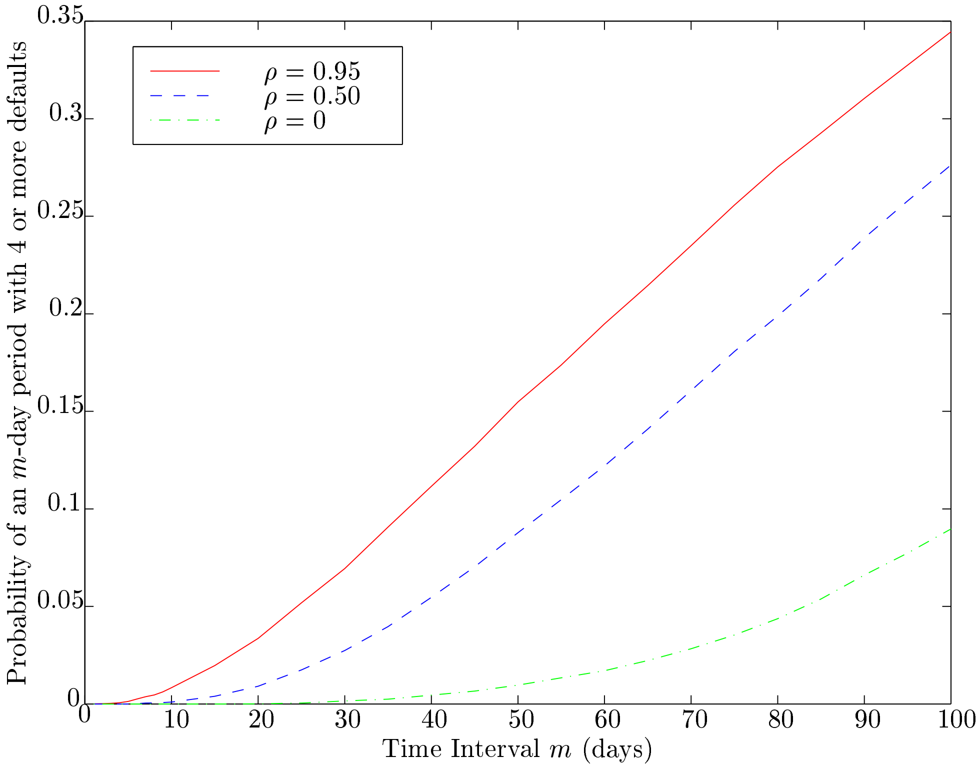

Figure 13: Probability of an m-day period within 10 years having 4 or more defaults (1,000 entities, intensity exponential Ornstein-Uhlenbeck, parameters $\theta = \ln(0.0017)$, $\sigma = 1$, $\kappa = 0.5$, pair-wise intensity shock correlation ρ).