Estimation of Physical Intensity Models for Default Risk

MICHEL DENAULT
GENEVIÈVE GAUTHIER
JEAN-GUY SIMONATO*

The estimation of physical intensity processes in the context of default risk is investigated here. Using data from Moody's Corporate Bond Default Database, a term structure of default probabilities for different rating classes is constructed each year from 1970 to 2001. Two specifications used for modeling the dynamics of the (risk-neutral) intensity process in the bond-pricing literature are then examined empirically: the Ornstein–Uhlenbeck and square-root cases. The results reveal that the Ornstein–Uhlenbeck case is not an adequate modeling alternative with a rejection of this specification in five out of seven credit classes and nonsignificant mean reverting behavior for all credit classes. The square-root

The authors thank K. Hammami for his generous help with the default data, as well as HEC Montréal for its financial support. Denault and Gauthier acknowledge the financial support of the National Science and Engineering Research Council of Canada (NSERC).

*Correspondence author, Service de l'enseignement de la finance, HEC Montréal, 3000 Côte-Sainte-Catherine, Montréal, Québec, Canada H3T 2A7. Tel: +1-514-340 6807, Fax: +1-514-340-5632, e-mail: jean-guy.simonato@hec.ca

Received October 2007; Accepted February 2008

- *Michel Denault is an Associate Professor at the Management Science Department of HEC Montréal, Montréal, Québec, Canada.*
- *Geneviève Gauthier is an Associate Professor at the Management Science Department of HEC Montréal, Montréal, Québec, Canada.*
- *Jean-Guy Simonato is a Professor at the Finance Department of HEC Montréal, Montréal, Québec, Canada.*

The Journal of Futures Markets, Vol. 29, No. 2, 95–113 (2009)
© 2008 Wiley Periodicals, Inc.
Published online in Wiley InterScience (www.interscience.wiley.com).
DOI: 10.1002/fut.20353

case obtains better results with four credit classes out of seven for which this specification cannot be rejected and significant mean reversion parameters in many cases. © 2008 Wiley Periodicals, Inc. Jrl Fut Mark 29:95–113, 2009

INTRODUCTION

Measuring, controlling and pricing credit risk, the risk associated with the possibility of default on a loan, are among the major challenges faced by financial institutions and bond portfolio managers. Many inputs or estimates are called for when performing such tasks. Among these, one of the most important is an estimate of the default probability and its dynamics through time.

In the financial literature, the approaches used for estimating creditworthiness and default probabilities can be classified into four categories. A first group estimates the default probability using the average frequency with which obligors of the same rating have defaulted. Agencies such as Moody's and Standard and Poor's collect data to perform these rating estimates. The main advantage of this approach is its simplicity and wide accessibility. A second approach uses statistical techniques and data such as balance sheets, current market conditions or past performance to estimate the probability that a firm will default. These statistical models include discriminant analysis and qualitative-response models as in Altman (1968) and Lo (1986) and duration models as in Shumway (2001). These methods depend on the availability of accurate and up-to-date financial statements and market condition data. A third group uses the structural bond-pricing approach started by Merton (1974). Using assumptions about the dynamics of the asset value of a firm and a default mechanism, a model allowing the computation of a default probability is obtained. Implementing such models relies on a mix of balance sheet and stock price data. An important advantage is the clear economic intuition about the default mechanism. It also uses up-to-date stock market data, which are easily available and contain the most recent information about events affecting the credit quality of the firm. The main drawback is that ad hoc simplifications about debt structures and maturities are often needed. Recent empirical applications in the literature are Brockman and Turtle (2003) and Ericsson and Reneby (2004). Finally, a fourth group consists of models from the reduced form approach, pioneered by Jarrow and Turnbull (1995) and Duffie and Singleton (1999). In these models, the instantaneous conditional default probability (hereafter called the intensity) is modeled directly with an assumption about its dynamics. This approach to the problem allows transparent and parsimonious forms for the default probability and its possible evolution though time. It can also be carried out using yield spread data alone and without some of the ad hoc assumptions often needed by the structural models. This research is about this fourth class of models. More precisely, the parameters' estimation of the intensity process driving the distribution of the default time is studied.

Several studies have focused directly or indirectly on estimating the intensity process of reduced form models (see Driessen, 2005; Duffee, 1999). These studies have relied on corporate bond price data and are thus impeded by several factors such as recovery rates, taxation and liquidity (see, for example, Driessen, 2005; Elton, Gruber, Agrawal, & Mann, 2001). In such contexts, bond prices are modeled as expectations of discounted payoffs under a risk-neutral measure, leading to an estimate of the intensity process under some risk-free pricing measure. Although this is useful for pricing purposes, it produces default probabilities under a measure that does not reflect the physical (true) probabilities of observing a default. In this study, the focus is on estimating intensity process parameters under the physical probability measure. The knowledge of the intensity process under the physical probability measure can be useful in applications such as value-at-risk estimation. In this case, the analyst needs to assess the real-world probabilities of defaulting and their dynamics for computing the estimates of the percentile of the portfolio distribution. Estimates of physical intensities are also interesting for pricing purposes. One could explore mappings and factors linking the risk-neutral and physical intensities. As mentioned in Duffie and Singleton (2003), this knowledge is interesting for the pricing of newly issued bonds or for discriminating between some of the models proposed in the literature.

The approach adopted here for estimating physical intensities dynamics uses data from Moody's Corporate Bond Default Database and proceeds in two steps. In a first step, for each year from 1970 to 2001, an estimate of the term structure of survival probabilities for a given rating class is obtained. The focus is on estimates of intensities for a whole class of firms grouped by credit ratings rather than for a single firm. The results will thus be interpreted as estimates for a typical or average firm within a rating class. It is then assumed that economic agents use rating transition data to form their anticipations about survival and default probabilities for various horizons. When forming their anticipations, it is also assumed that a one-year time window is used because the most recent data contain the most relevant information. Implementing the above assumptions is carried out with estimates of Markovian generators governing the credit migrations as in Lando and Skodeberg (2002). Each year, an estimate of this generator is computed with the transition and default data from the past year. A term structure estimate implied from the generator is then computed.

In a second step, relying on the term structures constructed in the previous step, two forms of intensity dynamics that have been used in the corporate bond-pricing literature are then examined: the Ornstein–Uhlenbeck case and the square-root case. With the theoretical relation between the survival probabilities at different time horizons and the postulated process, parameter and intensities estimates are obtained using a Kalman filter approach. The advantage of using the approach described above is that, unlike the procedures relying on corporate bond prices, ad hoc assumptions about recovery rates, taxes or liquidity are not required. The estimates of intensities are thus free of these considerations. The costs of avoiding these are the use of noisy estimates for the term structures of survival probabilities, which will contain approximation errors. The Kalman filter estimation approach used here is however well suited to handle such a case as it assumes that the observed survival probabilities contain random observation errors.

In related studies, Berndt, Douglas, Duffie, Ferguson, and Schranz (2005) and Berndt (2007) also look at the estimation of physical intensity models with estimated default probabilities. Berndt et al. (2005) develop a methodology for physical intensity estimation, although this is not their main purpose. Berndt (2007) on the other hand builds upon the methodology of Berndt et al. (2005) and tests several one-factor default intensity specifications for individual firms and firms grouped by industries. Her results support that a preferred specification is an Ornstein–Uhlenbeck for the logarithm of the default intensities, whereas the Ornstein–Uhlenbeck and square-root specifications for the level of the default intensities are rejected. Her approach differs from the one adopted here in two main aspects. First, the default probabilities estimates are obtained from Moody's KMV expected default frequency (EDF) measures. An advantage of using EDF data is the capacity of obtaining estimates for individual firms. This is not possible with the approach adopted here. EDF data however, just like the estimated probabilities used here, are not true measures of default probabilities and also suffer from censoring and missing observations. Another main difference between Berndt's (2007) approach and the one adopted here is that the potential misspecification of the model used to obtain the default probabilities is ignored in Berndt (2007). The approach used here handles misspecification by allowing measurement errors in estimated default probabilities.

The rest of the paper is organized as follows. The second section presents the model. The third section examines the estimation strategy and the fourth section describes the data and the empirical results. The fifth section concludes.

INTENSITY MODELS

Intensity models for default probabilities postulate that the distribution of the default time τ can be characterized by an intensity process. Conditional on the path of the intensity $\{\lambda_t : t \geq 0\}$, default occurs according to a Poisson arrival with this time-varying intensity. In such a setting, given that a firm has survived to the current time, t, the survival probability to a future time T conditional upon the information at time $t < T$ is

$$\Pr_t[\tau > T] = E_t\left[\exp\left(-\int_t^T \lambda_s\, ds\right)\right].$$

See, for example, Lando (2004) or Duffie and Singleton (2003). Various forms of intensity processes have been suggested in the literature. This study focuses on two well-known members of the exponential-affine family introduced in Duffie and Kan (1996) for the term structure of interest rates. More specifically, the one-dimensional Ornstein–Uhlenbeck process and square-root process are examined, which can be written as a mean reverting process

$$d\lambda_t = \kappa(\theta - \lambda_t)dt + \sigma\lambda_t^\delta dW_t \tag{1}$$

with κ the speed of mean reversion, θ the long run mean, σ the diffusion parameter, W a Brownian motion and $\delta = 0$ for the Ornstein–Uhlenbeck case and $\delta = \frac{1}{2}$ for the square-root case. With these specifications, an explicit solution for the survival probability can then be written as

$$\Pr_t[\tau > T] = S_t(T - t, \lambda_t) = A(\psi, T - t)\exp(-B(\psi, T - t)\lambda_t) \tag{2}$$

where $\psi = [\kappa, \theta, \sigma]$ and

$$A(\psi, u) = e^{[\gamma(B(\psi,u) - u) - \sigma^2 B^2(\psi,u)/4\kappa]}$$
$$B(\psi, u) = \frac{1}{\kappa}[1 - e^{-\kappa u}]$$

with $\gamma = \theta - (\sigma^2/2\kappa^2)$ for the Ornstein–Uhlenbeck case and

$$A(\psi, u) = \left[\frac{2\gamma e^{[(\kappa+\gamma)u]/2}}{(\kappa + \gamma)(e^{\gamma u} - 1) + 2\gamma}\right]^{2\kappa\theta/\sigma^2}$$

$$B(\psi, u) = \frac{2(e^{\gamma u} - 1)}{(\kappa + \gamma)(e^{\gamma u} - 1) + 2\gamma}$$

with $\gamma = \sqrt{\kappa^2 + 2\sigma^2}$ for the square-root case. Note that in this setting, the intensity process of the Ornstein–Uhlenbeck case may be negative, which contradicts one of the assumption. This specification is nevertheless sometimes used because of its convenience. See, for example, Schönbucher (2002).

The goal in this paper is to assess whether the theoretical relations provided by Equation (2) are adequate descriptions of physical survival probabilities. Estimates and inference for parameters k, θ and σ for the Ornstein–Uhlenbeck and square-root cases for groups of firms within a rating class. The authors will therefore get estimates for a "typical" or representative firm within a rating class will also be obtained. Therefore, the estimates correspond to a "typical" or representative firm within a rating class.

ESTIMATION APPROACH

To carry out the theoretical relation provided by the survival probability expression, a time series of observations for the term structure of survival probabilities are required. These term structures are not observed and need to be estimated. For this purpose, it is assumed that, within a year, economic agents observe the transition from one credit class to another and the defaults in the various classes. Using this information, they form their anticipations about the probability of defaulting for various time horizons in each credit classes. There is thus an assumption of time homogeneity made when constructing the term structure. It is also assumed that, each year, economic agents discard the old data and use data from the past year because this information is more representative of the current market conditions. To implement empirically the above description, estimates of Markovian generators governing the credit migrations as in Lando and Skodeberg (2002) are used. Each year, an estimate of this generator is computed with the transition and default data from the past year. A term structure estimate of the survival probabilities is then implied from the generator. A precise description of the methodology is given in section Data and Results.

With the data obtained as above, an approach based on the Kalman filter for estimating the parameters of the unobserved intensity is used. This recursive algorithm allows the estimation of the unobserved intensity process and the model's parameters. For background on the Kalman filter, see Harvey (1989). In the case where the intensity belongs to the Gaussian family (the Ornstein–Uhlenbeck case), the Kalman filter provides a maximum likelihood estimator for the parameters. For the square-root case, this approach provides an approximate quasi-maximum likelihood estimator. Results from Monte Carlo studies in Duan and Simonato (1999), De Jong (2000) and Duffee and Stanton (2004) show that this approach generates valid results for estimation and inference.

To implement the Kalman filter procedure, one needs to cast the model given by Equations (1) and (2) in a state space form, which consists of two equations: a measurement equation and a transition equation, whose form is described as follows. Let $Y_t(k) = (1/k)\ \ln S(k, \lambda)$ be a log-transformation of the survival probability with $k = 1, 2, \ldots, n$ representing the horizon of the survival probability in number of years. Because estimates for the term structure of survival probabilities rather than the true term structures of survival probabilities are used here, a measurement error is added to the log transform of the survival probabilities equation:

$$Y_t(k) = \frac{1}{k}\ln A(\psi, k) - \frac{1}{k}B(\psi, k)\lambda_t + \varepsilon_{t,k}.$$

At a given time t, the estimates of the survival probabilities for n different time horizons are obtained. The variables $Y_t(k)$, k $\{1, 2, \ldots, n\}$ are therefore stacked to obtain the following system:

$$\begin{pmatrix} Y_t(1) \\ \vdots \\ Y_t(n) \end{pmatrix} = \begin{pmatrix} \frac{1}{1}\ln A(\psi, 1) \\ \vdots \\ \frac{1}{n}\ln A(\psi, n) \end{pmatrix} - \begin{pmatrix} \frac{1}{1}B(\psi, 1) \\ \vdots \\ \frac{1}{n}B(\psi, n) \end{pmatrix}\lambda_t + \begin{pmatrix} \varepsilon_{t,1} \\ \vdots \\ \varepsilon_{t,n} \end{pmatrix} \tag{3}$$

where the sequence $\{\varepsilon_t = (\varepsilon_{t,1}, \ldots, \varepsilon_{t,n})^{\mathrm{T}} : t \in \{1, 2, 3, \ldots\}\}$ consists of independent and identically distributed $N(0, \Sigma_\varepsilon)$ random vectors. It is further assumed that Σ_ε is diagonal. The system given by Equation (3) is referred to as the measurement or observation equation in the context of the Kalman filter.

As discussed in Duan and Simonato (1999), the transition or state equation can be written using the conditional expected values and variance of the intensity because these are affine functions. Let $\mu(\lambda_t, \psi) = E[\lambda_{t+1} | \lambda_t]$ and $V(\lambda_t,\psi) = Var[\lambda_{t+1} | \lambda_t]$ be the conditional expectation and variance of the intensity. Then the transition equation is given by

$$\lambda_{t+1} = \mu(\lambda_t, \psi) + \sqrt{V(\lambda_t, \psi)}\eta_{t+1}$$

where $\{\eta_t : t = 1, 2, 3 \ldots\}$ is a sequence of zero mean and unit variance independent Gaussian noises. Note that the transition equation is exact in the case where the intensity is an Ornstein–Uhlenbeck process but is an approximation when the intensity is assumed to be a square-root process. The specific forms taken by $\mu(\cdot)$ and $V(\cdot)$ will depend on the process assumed for the dynamic of the intensity. For the Ornstein–Uhlenbeck case,

$$\mu(\lambda_t, \psi) = \theta(1 - e^{-\kappa h}) + e^{-\kappa h}\lambda_t$$

and

$$V(\lambda_t, \psi) = \frac{\sigma^2}{2\kappa}(1 - e^{-2\kappa h})$$

where h is the length of time in years between two time series observations. For the square-root process,

$$\mu(\lambda_t, \psi) = \theta(1 - e^{-\kappa h}) + e^{-\kappa h}\lambda_t$$

and

$$V(\lambda_t, \psi) = \lambda_t \frac{\sigma^2}{\kappa}(e^{-\kappa h} - e^{-2\kappa h}) + \theta\frac{\sigma^2}{2\kappa}(1 - e^{-\kappa h})^2.$$

To summarize, the approach to the estimation problem makes several assumptions. It is first assumed that the estimated term structures of default probabilities are the appropriate default probabilities used by market participants. This assumption is however not too constraining as measurement errors accounting for the possible misspecification of the estimated probabilities are introduced. Another source of errors is the use of an approximate quasi-maximum likelihood estimator for the square-root case with unknown theoretical properties. As noticed above, several Monte Carlo studies show the validity of this approach. It should also be noticed that in this case, the estimated intensities might become negative, which is violating a basic assumption, as the intensity process is restricted to positive values. In such cases, the ad hoc procedure of truncating these negative estimates to zero is used. It is difficult to assess how this procedure affects the estimates. The impact should however be modest if the number of truncated estimates is small. The results should be interpreted with this caveat in mind.

DATA AND RESULTS

A description of Moody's default database is first presented. Numerical results relying on the method described above are then provided.

Data

To get estimates of the generators each year, rating and transition data are processed as in Christensen, Hansen, and Lando (2004). The rating transition histories used to estimate the generator are taken from the January 2002 version of Moody's Corporate Bond Default Database. Only issuers domiciled in United States and having at least one senior unsecured estimated rating are considered. The database contains 5,719 issuers (in all industry groups) with 46,305 registered debt issues and 23,666 rating observations. For each issuer the number of default dates in the Master Default Table is checked. A total of 1,041 default dates for 943 issuers in the period 1970–2001 are obtained. Some issuers (91) had more than one default date. In the rating transition histories, there are 728 withdrawn ratings that are not the last observation of the issuer. These irrelevant withdrawals were eliminated.

When estimating the generator, the most important and difficult task is to get a proper definition of default. Default dates are treated here as they are in Christensen et al. (2004). First, all the nonwithdrawn-rating observations up to the date of default have typically been unchanged. However, ratings occuring within a week before the default date were eliminated. Rating changes observed after the date of default were eliminated unless the new rating reached the B3 level or higher and the new ratings were related to debt issued after the date of default. Such cases were treated as related to a new issuer. It is important to emphasize that the first rating date of the new issuer is the latest date between the date of the first issue after default and the first date we observe an issuer rating higher than or equal to B3. The same treatment is applied for the case of two and three default dates. Finally, few issuers have a registered default date before the first rating observation in the Senior Unsecured Estimated Rating Table. In these cases, it was considered that there was no default. With this procedure, the revised database ends up with 5,821 issuers with 965 default dates.

Aggregating all rating notches yields the nine usual ratings, Aaa, Aa, A, Baa, Ba, B, Caa-C, Default and NR (Not Rated), with 15,564 rating observations. When estimating the survival probabilities, eight different credit ratings were considered: Aaa, Aa, a, Baa, Ba, B, Caa and Default. For a given year t, the 8×8 generator of the Markov chain describing the rating migrations is estimated with the approach proposed in Lando and Skodeberg (2002). This estimate is obtained with the rating migrations for this specific year. The estimate of the transition matrix $\mathbf{M}_{t,k}$ for a time horizon of k years may be obtained from the generator, $\mathbf{G}_t$, using

$$\mathbf{M}_{t,k} = \exp(k\mathbf{G}_t) = \sum_{i=0}^{\infty} \frac{(k\mathbf{G}_t)^i}{i!}.$$

The survival probabilities are then computed for a horizon of $k = \{1, 2, 3, 4\}$ years for the eight different credit ratings as $\mathbf{e} - \mathbf{m}_{t,k}$, where $\mathbf{e}$ is an 8×1 vector of ones and $\mathbf{m}_{t,k}$ is the last column of $\mathbf{M}_{t,k}$.

Table I presents summary statistics regarding the computed default probabilities obtained from the survival probabilities for the different rating classes. In some rating classes, the estimated probabilities are zero. A detailed look at the data reveals that these cases are rather infrequent (8 instances out of 224 cross-categories for rating and year) and are found in the following ratings and years: Aaa for years 1973, 1975, 1978 and 1979; Aa for the year 1978; Caa for years 1978, 1979 and 1981. In principle, a zero default probability should not be

TABLE I
Summary Statistics on Annual Observations of Estimated Default Probabilities for Years 1970–2001

	1 Year	*2 Years*	*3 Years*	*4 Years*
		Aaa		
Mean	0.00001167	0.00005457	0.00013960	0.00027625
Median	0.00000002	0.00000040	0.00000237	0.00000839
Std.	0.00003871	0.00017594	0.00043716	0.00083923
Max	0.00019594	0.00087233	0.00212623	0.00401613
Min	0.000e+000	0.000e+000	0.000e+000	0.000e+000
		Aa		
Mean	0.00003456	0.00014915	0.00036172	0.00069283
Median	0.00000141	0.00001251	0.00004395	0.00010719
Std.	0.00009644	0.00038830	0.00087878	0.00157387
Max	0.00047271	0.00188868	0.00424988	0.00756704
Min	0.000e+000	0.000e+000	0.000e+000	0.000e+000
		A		
Mean	0.00016401	0.00056505	0.00127811	0.00235794
Median	0.00003020	0.00014214	0.00036560	0.00067842
Std.	0.00042875	0.00106976	0.00212595	0.00367432
Max	0.00231594	0.00484851	0.00829725	0.01501596
Min	1.634e−007	2.207e−006	7.551e−006	1.811e−005
		Baa		
Mean	0.00150927	0.00416045	0.00798348	0.01293471
Median	0.00037180	0.00137877	0.00293487	0.00514017
Std.	0.00203813	0.00501036	0.00916558	0.01448697
Max	0.00735245	0.01733422	0.03077074	0.04943764
Min	1.384e−005	1.095e−004	3.634e−004	7.762e−004
		Ba		
Mean	0.01064171	0.02595188	0.04465796	0.06548246
Median	0.00834628	0.01937027	0.03104777	0.04187505
Std.	0.00992224	0.02248155	0.03769198	0.05455828
Max	0.04155428	0.08011173	0.13221063	0.19686810
Min	4.150e−004	1.591e−003	3.431e−003	5.850e−003
		B		
Mean	0.07056594	0.13756550	0.19774688	0.25064015
Median	0.06504085	0.12475919	0.18929791	0.23416089
Std.	0.04906585	0.08696072	0.11671285	0.14018393
Max	0.23058479	0.38453545	0.48863711	0.56021486
Min	6.486e−005	2.458e−004	5.245e−004	8.853e−004
		Caa		
Mean	0.29268655	0.45140739	0.54635929	0.60795588
Median	0.27206638	0.45945702	0.59363006	0.65673053
Std.	0.18097856	0.24552669	0.27158200	0.28377459
Max	0.62911009	0.86244067	0.94898063	0.98107743
Min	0.000e+000	0.000e+000	0.000e+000	0.000e+000

Note. Columns 2–5 present the summary statistics for default probabilities with horizons of 1–4 years. The default probabilities at a given year have been obtained by estimating the generator of the Markov chain of credit rating migrations with the Lando and Skodeberg (2002) approach described in subsection data and the previous year of default data from Moody's (2002) database.

possible if successive downgrades to a class with a positive default probability is possible. Because of this, these zero default probabilities will be treated as missing observations during the estimation procedure. In the context of the Kalman filter framework, missing observations can be handled easily with the procedure described in section 3.4.7 of Harvey (1989). The results for these rating classes will thus have to be interpreted with this caveat in mind. For all rating classes, the average default probabilities are increasing with the horizon and decreasing with credit quality. In general, the first 10 years of the sample show small default probabilities that increase in the second 10 years' portion (years 1980–1990) to decrease again for the rest of the sample.

Estimation Results for the Ornstein–Uhlenbeck Case

Table II presents the estimation results for the Ornstein–Uhlenbeck case implemented with the data generated with the above procedure. The table reports, for each rating class, the parameter estimates with the estimated standard deviation in parenthesis. The estimates for θ, the long run average, are increasing as credit quality is decreasing. This is expected and implies that the unconditional instantaneous probability of default is increasing as the credit quality is decreasing. Generally, this parameter is not estimated with much precision.

TABLE II

Estimation Results for the Ornstein–Uhlenbeck Specification with Annual Observations on 1, 2, 3 and 4 Years' Survival Probability Series for Years 1970–2001

	Rating						
Parameter	*Aaa*	*Aa*	*A*	*Baa*	*Ba*	*B*	*Caa*
θ	0.000152	0.000620	0.006495	0.037034	0.152480	0.197650	0.515101
–	(0.000208)	(0.000546)	(0.003192)	(0.015395)	(0.060396)	(0.075946)	(0.125929)
κ	0.066657	0.056507	0.035612	0.029695	0.030531	0.055140	0.124365
–	(0.082917)	(0.039423)	(0.014453)	(0.011738)	(0.012185)	(0.024181)	(0.052327)
σ	0.000060	0.000143	0.000845	0.003461	0.014392	0.060194	0.221705
–	(0.000028)	(0.000051)	(0.000180)	(0.000462)	(0.002687)	(0.009843)	(0.032946)
σ_{ε_1}	3.93e−008	7.31e−008	0.000198	0.000708	0.003089	0.013555	0.071955
–	(0.007806)	(0.016809)	(0.000071)	(0.000148)	(0.000592)	(0.003314)	(0.013730)
σ_{ε_2}	0.000051	0.000102	1.33e−011	4.85e−011	5.97e−009	0.005726	0.033938
–	(0.000021)	(0.000034)	(232.3789)	(602.9217)	(21.0965)	(0.002064)	(0.005950)
σ_{ε_3}	0.000111	0.000207	0.000223	0.000705	0.002687	1.65e−006	0.000016
–	(0.000040)	(0.000067)	(0.000068)	(0.000135)	(0.000501)	(3.361877)	(2.128010)
σ_{ε_4}	0.000178	0.000316	0.000468	0.001399	0.004948	0.004404	0.030345
–	(0.000063)	(0.000101)	(0.000121)	(0.000223)	(0.000849)	(0.001947)	(0.004600)
Chi-square	18.56	14.98	21.41	14.34	18.85	18.95	23.33
P-value	0.0050	0.0204	0.0015	0.0260	0.0044	0.0042	0.0007

Note. The table reports the point estimates for the parameters with the estimated standard errors in parenthesis. Chi-square and *P*-value are the results of a Lagrange multiplier test with six degrees of freedom examining the null hypothesis of zero values for additional parameters in the measurement equation.

At the 1% confidence level, the null hypothesis of a zero-valued parameter cannot be rejected for classes from Aaa to Ba. The mean reversion parameter κ is estimated to be small for all credit classes and not statistically different from zero at the 1% level for all credit classes. The standard deviation parameter of the Ornstein–Uhlenbeck process, σ, is significant in all cases except Aaa and its magnitude also increases as credit quality decreases. The measurement error's standard deviation estimates of Equation (3) are significantly different from zero in most cases with small estimated values when compared with the standard deviation parameter of the Ornstein–Uhlenbeck process. Interestingly, for all credit classes, one of the maturities shows a very small measurement error standard deviation estimate. This implies that the unobserved intensities and the time series data associated with this maturity are similar.

The last two lines of the table report the results of a Lagrange multiplier test examining whether the theoretical restrictions prescribed in the measurement equations are validated by the data. Specifically, this test examines whether additional free parameters placed in the measurement equation are statistically different from zero. For details about this test, see Duan and Simonato (1999). The results show that a null hypothesis of zero values for these additional free parameters cannot be rejected at the 1% level for Aa and Baa classes only.

The estimated time series for the unobserved intensities are plotted in Figure 1. For the Aaa, Aa, A and Baa classes, a minority of the estimated numbers are positives with 9, 7, 9 and 14 positive estimates out of 32 for each rating class. For the Ba, B and Caa classes, the number of positive estimates is 24, 30 and 31. The graphs show increases in the estimated intensities in the middle portion of the sample. The high intensity periods are different for some categories. For example, the firms rated A experienced a decrease in estimated intensities around 1990, whereas an increase was estimated for this period for firms rated B.

Figure 2 shows the forward default rate curves implied by the estimated parameters for each credit classes. These forward default rate curves can be interpreted as the rate of default arrival at a given time, T, conditional only on survival up to time T.[1] As discussed in Duffie and Singleton (2003), the shape of these curves, given $\lambda_0 = \hat{\theta}$, depends on the relative values of the volatility and mean reversion parameters. Higher values of σ tend to lower the probability of default given survival as time increases because of the convexity of the survival probability function. On the other hand, higher values of the mean reversion parameter κ tend to keep the forward default rate close to the initial

[1]The theoretical relation linking forward default rates and survival probabilities is $S_0(T, \lambda_0) = \exp(-\int_0^T f_t\, dt)$. The forward default rates shown in the graphs are computed recursively with the discrete approximation $S_0(T, \lambda_0) = \exp(-\sum_{i=0}^{n-1} \tilde{f}_i \Delta t)$ with $\Delta t = 1$, $n = T/\Delta t$ and $\tilde{f}_0 = \lambda_0$.

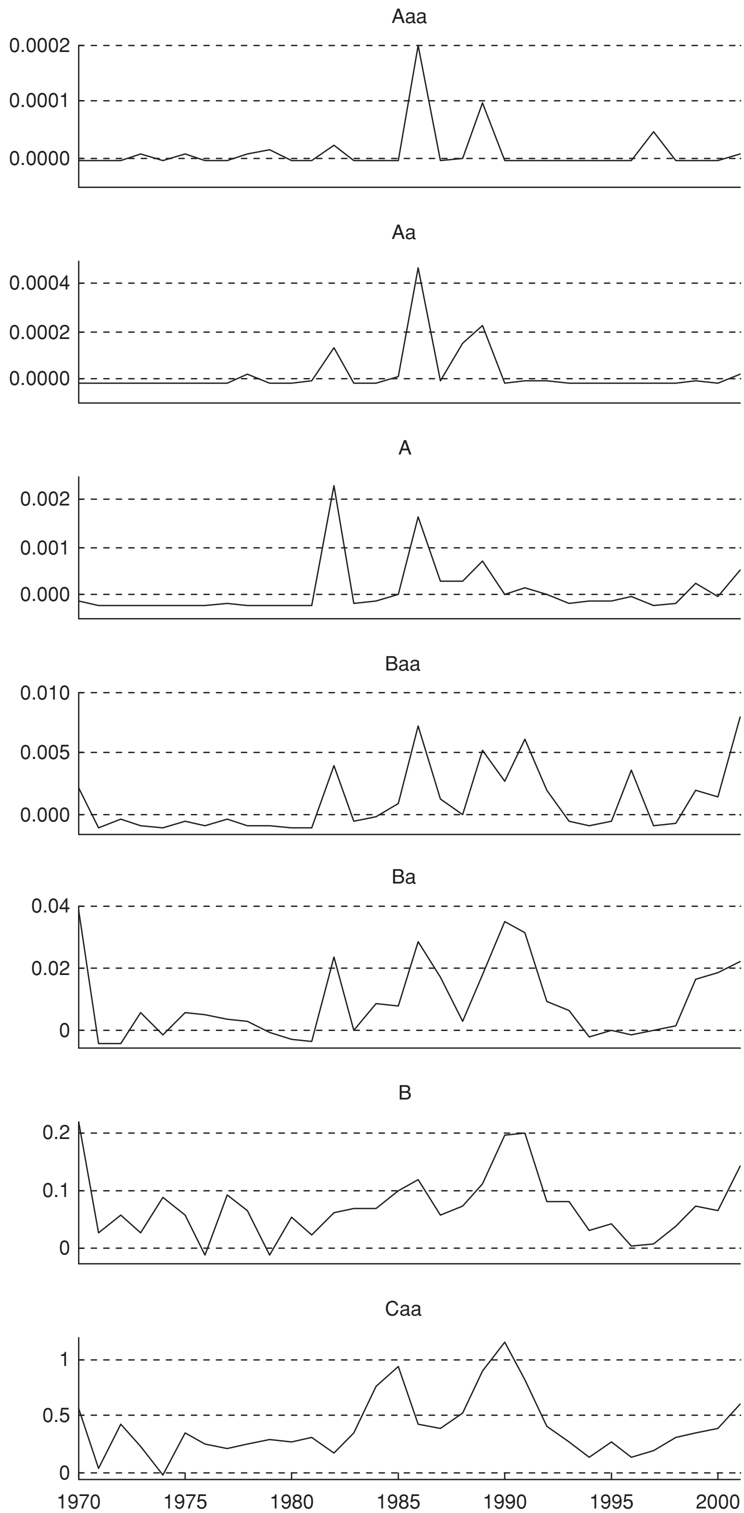

FIGURE 1
Estimated intensities for the Ornstein–Uhlenbeck specification.

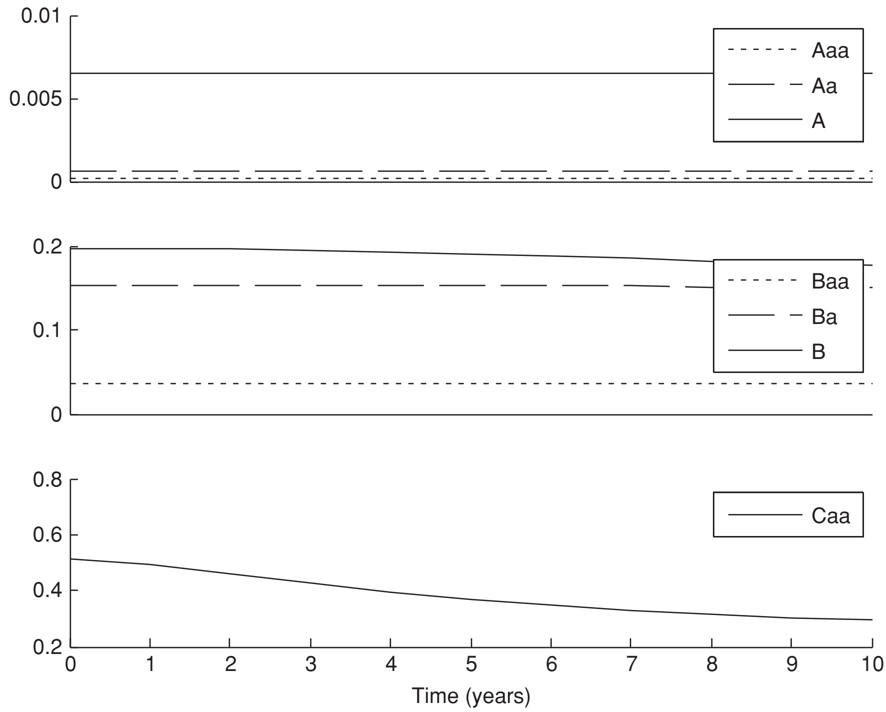

FIGURE 2
Forward default rates for the Ornstein–Uhlenbeck case.

default rate. There is thus a trade-off between the effects of the volatility and mean reversion parameters. For the Aaa, Aa, A, Baa and Ba credit classes, these curves are flat showing the effect of the small estimated volatility parameter relative to the mean reversion parameter. For B and Caa, these curves are downward sloping with the estimated volatility parameter now larger than the mean reversion parameter. These downward-sloping curves imply that firms in these classes will be expected to improve in credit quality (lower default rate), conditional on survival, as the horizon increases.

Estimation Results for the Square-Root Case

Table III presents the results for the square-root case. The estimated parameters are qualitatively similar to those of the Ornstein–Uhlenbeck case as the magnitude of all estimated parameters is generally increasing as credit quality is decreasing. When compared with the parameter estimates of the Ornstein–Uhlenbeck, the estimates obtained for the long run mean are generally higher but not different from zero at the 1% level for the Aaa, Aa, B and Caa classes. For the same categories, the mean reversion parameters are not

TABLE III

Estimation Results for the Square-Root Specification with Annual Observations on 1, 2, 3 and 4 Years' Survival Probability Series for Years 1970–2001

	Rating						
Parameter	*Aaa*	*Aa*	*A*	*Baa*	*Ba*	*B*	*Caa*
θ	0.000438	0.002844	0.030471	0.146541	0.451005	0.423428	1.317985
–	(0.000208)	(0.001135)	(0.005054)	(0.054421)	(0.164006)	(0.262894)	(0.880626)
κ	0.034894	0.007297	0.007419	0.006673	0.010510	0.028255	0.037647
–	(0.017917)	(0.003577)	(0.001742)	(0.001650)	(0.003313)	(0.013280)	(0.019001)
σ	0.014923	0.034145	0.058952	0.078275	0.121591	0.243385	0.423175
–	(0.005561)	(0.009057)	(0.009537)	(0.004571)	(0.010586)	(0.024985)	(0.056625)
σ_{ε_1}	2.70e−008	0.000208	0.000412	0.000696	0.003039	0.014285	0.039725
–	(0.011155)	(0.000065)	(0.000122)	(0.000098)	(0.000740)	(0.002760)	(0.007562)
σ_{ε_2}	0.000050	0.000106	0.000219	2.86e−007	6.14e−006	0.006177	0.000009
–	(0.000024)	(0.000033)	(0.000036)	(0.007968)	(0.116079)	(0.001087)	(0.786919)
σ_{ε_3}	1.09e−004	6.00e−008	1.00e−007	0.000697	0.002698	0.000002	0.034854
–	(0.000041)	(0.000936)	(0.004503)	(0.000113)	(0.000656)	(0.038834)	(0.005369)
σ_{ε_4}	0.000175	0.000111	0.000243	0.001387	0.005014	0.004910	0.064837
–	(0.000062)	(0.000033)	(0.000032)	(0.000174)	(0.000875)	(0.000734)	(0.010441)
Chi-square	25.49	15.41	26.22	15.72	17.55	16.84	15.80
P-value	0.0003	0.0173	0.0002	0.0153	0.0075	0.0099	0.0149

Note. The table reports the point estimates for the parameters with the estimated standard errors in parenthesis. Chi-square and *P*-value are the results of a Lagrange multiplier test with six degrees of freedom examining the null hypothesis of zero values for additional parameters in the measurement equation.

significant at the 1% level. These parameter estimates are smaller in magnitude than those obtained for the Ornstein–Uhlenbeck case. As discussed by Duan and Simonato (1999) and Duffee and Stanton (2004), these estimation results should be interpreted with care as a small estimated mean reversion parameter suggests a nearly integrated process. In such a case, the resulting distribution for the test statistic might differ from the approximate asymptotic normal distribution. These small parameter estimates also imply that the estimated factor can reach zero with a positive probability as $\hat{\sigma}^2 > 2\hat{\kappa}\hat{\theta}$ (see Cox, Ingersoll, & Ross, 1985). The magnitudes of the volatility parameter estimates are also larger than those obtained for the Ornstein–Uhlenbeck case and are all significantly different from zero at the 1% level.

The results of the Lagrange multiplier test show that the null hypothesis of zero values for additional parameters in the measurement equation is not rejected at the 1% level for the Aa, Baa, B and Caa categories. For Aaa and A, the restrictions are strongly rejected, whereas it is rejected with a *P*-value of 0.0075 for the Ba category.

The estimated time series for the unobserved intensities are plotted in Figure 3. Note that to avoid numerical difficulties associated with negative estimates of the intensity process during the Kalman filter recursions, the estimated

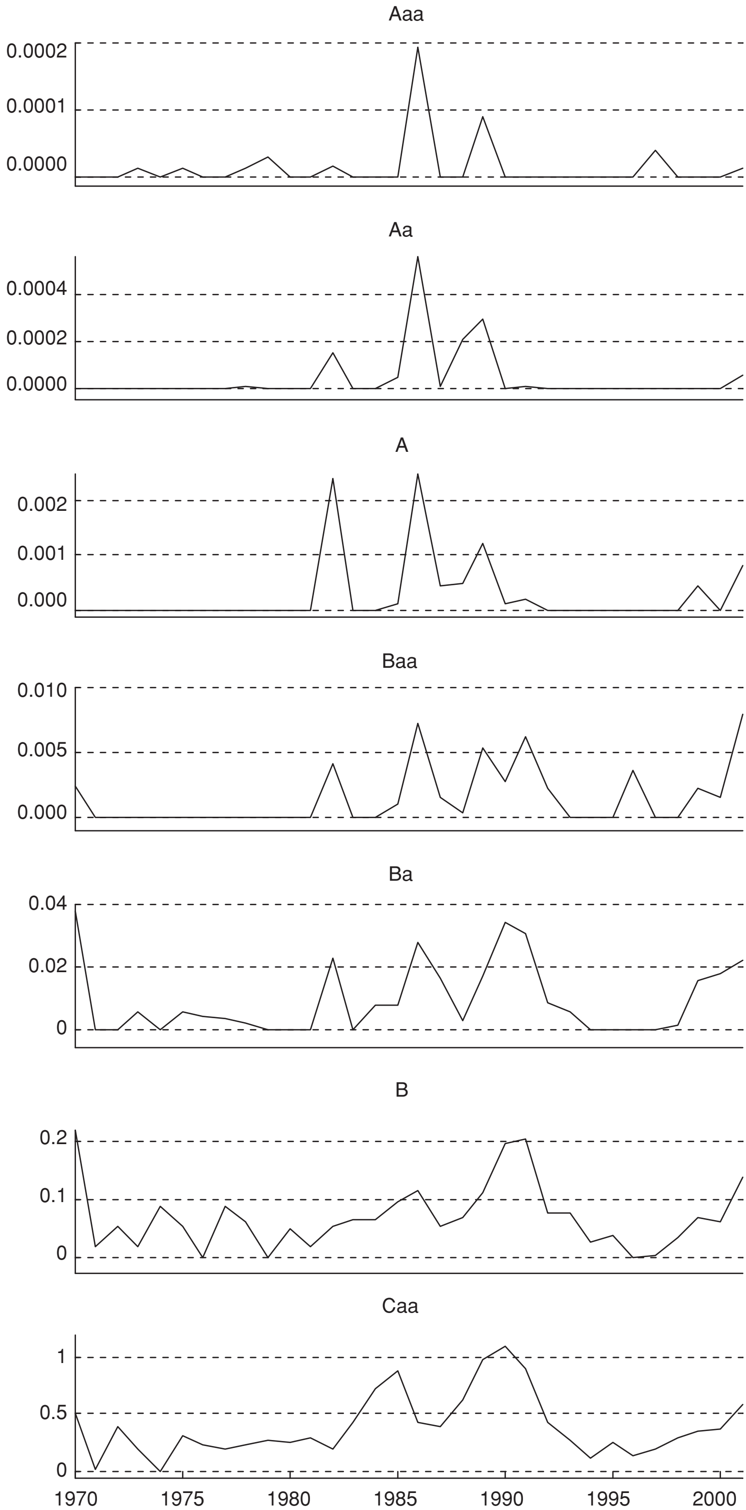

FIGURE 3
Estimated intensities for the square-root specification.

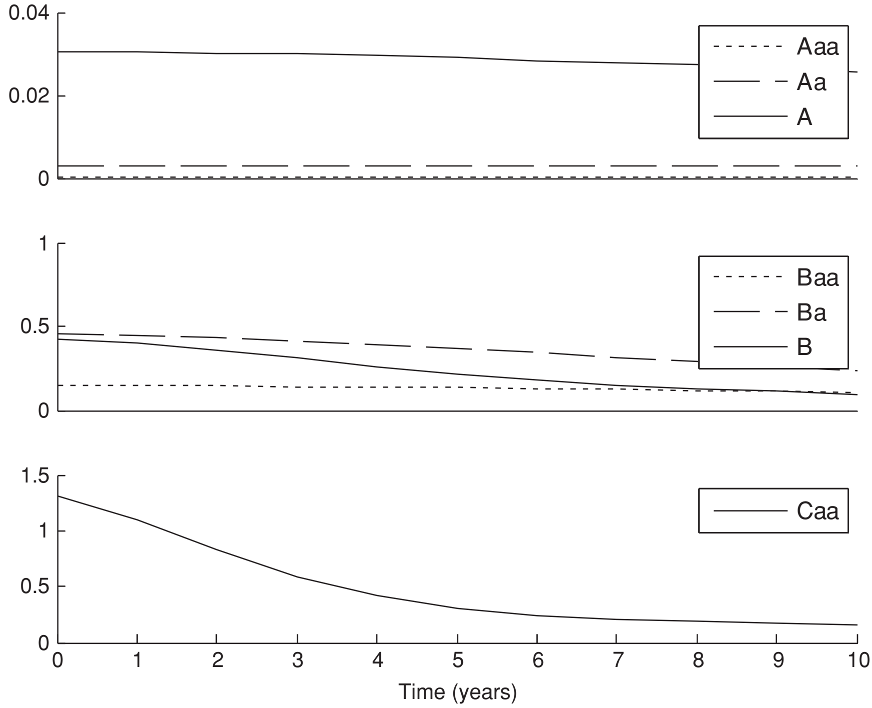

FIGURE 4
Forward default rates for the square-root case.

intensities were fixed to zero whenever a negative estimate was encountered. For Aaa, Aa, A and Baa classes, the majority of the estimated intensities are at the imposed lower bound of zero with 9, 12, 10 and 14 intensities being estimated greater than zero on a total of 32 estimated values for each rating class. For the Ba, B and Caa classes, the number of positive estimates for the intensities is 23, 29 and 31. Generally, the estimated intensities show similar patterns to those obtained for the Ornstein–Uhlenbeck case.

Figure 4 shows the forward default rate curves implied by the estimated parameters for each credit classes. The patterns are slightly different from the Ornstein–Uhlenbeck because of the smaller mean reversion and higher volatility parameters. For the high credit classes such as Aaa and Aa, these curves are flat, showing that the credit quality does not improve conditional upon survival. For A and Baa the curves are slightly downward sloping, whereas the Ba, B and Caa cases show clear negative slopes. These downward curves imply that a firm in these classes will be expected to improve in credit quality, conditional on survival, as the horizon increases. Finally, it should be noticed that the Caa class shows probabilities greater than one because of the very large estimated long run mean. This parameter was, however, not found to be statistically significant.

CONCLUSION

The estimation of physical intensity processes in the context of default risk was examined here. Unlike the previous literature that has relied on corporate bond price data to obtain these estimates, rating transition and default data from Moody's Corporate Bond Default Database were used. The estimates obtained are thus free of the usual assumptions about taxes, recovery rates and liquidity that are required when dealing with corporate bond data. Two specifications, both commonly used for the intensity dynamics modeling, have been examined: the Ornstein–Uhlenbeck and square-root models. Because the estimates are under the physical probabilities, they could be used in, for example, risk management applications where the dynamics of actual default probabilities are often required at aggregated credit level classes.

For the Ornstein–Uhlenbeck specification, the results of the robust Lagrange multiplier test are in general negative. At the 1% confidence level, the theoretical model for two credit classes out of seven cannot be rejected. For all credit classes, a zero value for the mean reversion parameter cannot be rejected. For the square-root case, the results are mixed but more supportive of the model as four credit classes out of seven cannot reject the robust Lagrange multiplier test. Mean reversion is also more present as three classes out of seven show significant mean reversion and long run mean parameters at the 1% level. These results indicate that this process often found in the corporate bond-pricing literature could be an adequate modeling choice for the riskier bond classes at the aggregate level.

The results should be interpreted with the limitations of the methodology in mind. A first limitation of their study is the assumption that a Markovian generator can be used to produce the estimates of the term structure each year. The term structures obtained this way can only be regarded as approximations of the unobserved true term structure of survival probabilities. Finally, it was implicitly assumed that firms have a certain degree of homogeneity within a rating category. A rejection of the specifications examined at the aggregate level does not necessarily invalidate the models at the firm's level.

REFERENCES

Altman, E. (1968). Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. Journal of Finance, 23, 589–609.

Berndt, A. (2007). Specification analysis of reduced-form credit risk models. Working paper, Tepper School of Business.

Berndt, A., Douglas, R., Duffie, D., Ferguson, M., & Schranz, D. (2005). Measuring default risk premia from default swap rates and EDFs. Working paper, Tepper School of Business.

Brockman, P., & Turtle, H. (2003). A barrier option framework for corporate security valuation. Journal of Financial Economics, 67, 511–529.

Christensen, J., Hansen, E., & Lando, D. (2004). Confidence sets for continuous-time rating transition probabilities. Journal of Banking and Finance, 28, 2575–2602.

Cox, J., Ingersoll, J., & Ross, S. (1985). A theory of the term structure of interest rates. Econometrica, 53, 385–407.

De Jong, F. (2000). Time series and cross-section information in affine term-structure models. Journal of Business and Economic Statistics, 18, 300–314.

Driessen, J. (2005). Is default event risk priced in corporate bonds? Review of Financial Studies, 18, 165–195.

Duan, J.-C., & Simonato, J.-G. (1999). Estimating and testing exponential-affine tem structure models by Kalman filter. Review of Quantitative Finance and Accounting, 13, 111–135.

Duffee, G. (1999). Estimating the price of default risk. Review of Financial Studies, 12, 197–226.

Duffee, G., & Stanton, R. (2004). Estimation of dynamic term structure models. Working paper, Hass School of Business.

Duffie, D., & Kan, R. (1996). A yield-factor model of interest rates. Mathematical Finance, 6, 379–406.

Duffie, D., & Singleton, K. (1999). Modeling term structures of defaultable bonds. The Review of Financial Studies, 12, 687–720.

Duffie, D., & Singleton, K. (2003). Credit risk: Pricing, measurement and management. Princeton: Princeton University Press.

Elton, E., Gruber, M., Agrawal, D., & Mann, C. (2001). Explaining the rate spread on corporate bonds. The Journal of Finance, 56, 247–277.

Ericsson, J., & Reneby, J. (2004). Implementing structural credit risk models using both stock and bond prices—an empirical study. Journal of Fixed Income, 13, 38–49.

Harvey, A. (1989). Forecasting structural time series model and the Kalman filter. Cambridge: Cambridge University Press.

Jarrow, R., & Turnbull, S. (1995). Pricing derivatives on financial securities subject to credit risk. Journal of Finance, 1, 53–85.

Lando, D. (2004). Credit risk modeling. Princeton: Princeton University Press.

Lando, D., & Skodeberg, T. (2002). Analyzing rating transitions and rating drift with continuous observations. Journal of Banking and Finance, 26, 423–444.

Lo, A. (1986). Logit versus discriminant analysis: A specification test and application to corporate bankruptcies. Journal of Econometrics, 31, 151–178.

Merton, R. (1974). On the pricing of corporate debt: the Risk structure of interest rates. Journal of Finance, 28, 449–470.

Schönbucher, P. (2002). A tree implementation of a credit spread model for credit derivatives. Journal of Computational Finance, 6, 1–38.

Shumway, T. (2001). Forecasting bankruptcy more accurately: A simple hazard model. Journal of Business, 74, 101–124.