Estimating the Price of Default Risk

Gregory R. Duffee
Federal Reserve Board

A firm's instantaneous probability of default is modeled as a translated square-root diffusion process modified to allow the process to be correlated with default-free interest rates. The parameters of the process are estimated for 161 firms. An extended Kalman filter approach is used that incorporates both the time-series and cross-sectional (term structure) properties of the individual firms' bond prices. The model is reasonably successful at fitting corporate bond yields, while key features of the term structures of yield spreads are captured in the signs and magnitudes of the resulting parameter estimates.

The risk of default affects virtually every financial contract. Therefore the pricing of default risk has received much attention; both from traders, who have a strong interest in pricing transactions accurately, and from financial economists, who have much to learn from the way such risks are priced in markets. The standard theoretical paradigm for modeling credit risks is the contingent claims approach pioneered by Black and Scholes. Much of the literature follows Merton (1974) by explicitly linking the risk of a firm's default to the variability in the firm's asset value. Although this line of research has proven very useful in addressing the qualitatively important aspects of pricing credit risks, it has been less successful in practical applications.[1] The lack of success owes to the difficulty of modeling realistic boundary conditions. These boundaries include both the conditions under which default occurs, and in the event of default, the division of the value of the firm among claimants. Firms' capital structures are typically quite complex and priority rules are often violated.

In response to these difficulties, an alternative modeling approach has been pursued in a number of recent articles, including Madan and Unal (1994), Duffie and Singleton (1995, 1997), and Jarrow and Turnbull (1995).[2]

I thank Torben Andersen, Mark Fisher, Dilip Madan, Ken Singleton, and Chunsheng Zhou for helpful conversations, and Kaushik Amin, Ravi Jagannathan (the editor), Francis Longstaff, two anonymous referees, and seminar participants at the 1997 Western Finance Meetings, Ohio State, and an NBER Derivatives Project meeting for helpful comments. The analysis and conclusions of this article are those of the author and do not indicate concurrence by other members of the research staff, by the Board of Governors, or by the Federal Reserve Banks. Address correspondence to Greg Duffee, Mail Stop 91, Federal Reserve Board, Washington, DC 20551, or email: gduffee@frb.gov.

[1] One illustration of this lack of success is that very few articles have even attempted to price particular instruments. The standard reference is Jones et al. (1984), who find that, even for firms with very simple capital structures, a Merton-type model is unable to price investment-grade corporate bonds better than a naive model that assumes no risk of default.

[2] Researchers also continue to extend Merton-style models to address these problems. Recent examples are Longstaff and Schwartz (1995) and Zhou (1997).

The Review of Financial Studies Spring 1999 Vol. 12, No. 1, pp. 197–226

At each instant there is some probability that a firm defaults on its obligations. Both this probability and the recovery rate in the event of default may vary stochastically through time. The stochastic processes determine the price of credit risk. Although these processes are not formally linked to the firm's asset value, there is presumably some underlying relation, thus Duffie and Singleton describe this alternative approach as a reduced form model.

This class of models is much more tractable mathematically than Merton-type models, raising hopes that it will be useful in practical applications. However, empirical implementation of this type of model is in its infancy. The few analyses that have been done to date fit such models to aggregate yield indexes, hence the ability of these models to price accurately the credit risks associated with particular instruments is unknown.

This article is an effort to determine, on a broad scale, whether one such model can describe the behavior of individual corporate bond prices successfully. The data consist of month-end observations on noncallable corporate bonds from January 1985 through December 1995; in total, more than 40,000 bond-price observations across 161 firms. The vast majority of the bonds are rated investment grade.

I follow Pearson and Sun (1994) by fitting the default-free term structure to a translated two-factor square-root diffusion model (although unlike Pearson and Sun, I estimate the model using the extended Kalman filter); *translated* means that a constant term is included. The model is extended to noncallable corporate bonds by assuming that the instantaneous probability that a given firm defaults on its obligated bond payments follows a translated single-factor square-root diffusion process, with a modification that allows the default process to be correlated with the factors driving the default-free term structure. Realistically there are a number of factors other than default risk that drive a wedge between corporate and Treasury bond prices, such as liquidity differences, state taxes, and special repo rates. Here all of these factors are subsumed into a stochastic process called a "default risk" process.

There are two goals to this exercise. The first is to determine how well such a model describes corporate bond yields, and in particular, to determine which features of the data are well-described by this model and which are not. The second is to see what the parameter estimates tell us about the behavior of individual firms' bond yields. Indexes of corporate bonds have been well-studied, but relatively little is known about the properties of individual firms' corporate bond yields, especially their time-series behavior.

On average the model fits corporate bond yields well. For the typical firm, the root mean squared error in yields is less than 10 basis points. However, not surprisingly, the model has difficulty matching certain features of the data. For example, there appear to be persistent fluctuations in the volatilities of yields (GARCH-like effects) that are not captured by the model. Perhaps more importantly, the parameter estimates for a typical firm rated, say, Baa, are different in important ways from the corresponding estimates for a typical firm rated Aa. This suggests parameter instability as firms' credit qualities change.

The parameter estimates reveal some important firm-level features of bond yields. For the typical firm, default risk is mean reverting under the true (physical) measure, but mean averting (i.e., nonstationary) under the equivalent martingale measure. Default risk is negatively correlated with default-free interest rates, although the strength of this negative correlation is weaker than reported elsewhere. In addition, for the typical firm, the instantaneous risk of default has a lower bound that exceeds zero. In other words, even if a firm's financial health dramatically improves, the model implies that yield spreads on the firm's bonds remain positive.

This article is not the first to examine empirically the pricing of credit risk using models of instantaneous default risk, although its use of individual bond data and the specification questions it addresses are new. Nielsen and Ronn (1996) analyze both corporate bond yield indexes and indexes of swap spreads. Indexes of midmarket swap spreads are also studied by Grinblatt (1995) and Duffie and Singleton (1997), although Grinblatt argues that this relation is driven by a liquidity yield to Treasury securities instead of default risk. Madan and Unal (1994) estimate a model of yields on certificates of deposit issued by roughly 300 thrift institutions.

Section 1 describes the model. The data are discussed in the Section 2 and the estimation procedure is discussed in Section 3. Section 4 reports the basic results, while Section 5 disaggregates the results by credit rating. Section 6 contains some specification tests of the model. Section 7 concludes the article.

1. A Model of Corporate Bond Prices

This model is adapted from Duffie and Singleton (1995, 1997), although Lando (1994) and Madan and Unal (1994) have similar frameworks. First consider default-free interest rates. The instantaneous nominal default-free interest rate is denoted r_t. Denote the time t price of a default-free bond that pays off a dollar at time T as $P(t, T, 0)$. (The third argument is the bond's coupon payment, which is zero.) The price of this bond is given by the expectation, under the equivalent martingale measure, of the cumulative discount rate between t and T:

$$P(t, T, 0) = E_t^Q \left\{ \exp\left[-\int_t^T r_u du \right] \right\}. \tag{1}$$

I follow Pearson and Sun (1994) by assuming that r_t equals the sum of a constant and two factors, $s_{1,t}$ and $s_{2,t}$, that follow independent square-root stochastic processes.

$$r_t = \alpha_r + s_{1,t} + s_{2,t}; \tag{2}$$

$$ds_{i,t} = \kappa_i(\theta_i - s_{i,t})dt + \sigma_i\sqrt{s_{i,t}}dZ_{i,t}, \quad dZ_{1,t}, dZ_{2,t} \text{ independent.} \tag{3a}$$

Under the equivalent martingale measure, these processes can be represented as

$$ds_{i,t} = (\kappa_i\theta_i - (\kappa_i + \lambda_i)s_{i,t})dt + \sigma_i\sqrt{s_{i,t}}d\hat{Z}_{i,t}, \quad d\hat{Z}_{1,t}, d\hat{Z}_{2,t} \text{ independent.} \tag{3b}$$

As is well known, the combination of Equations (2) and (3b) imply a closed-form solution for zero-coupon bond prices in which the bonds' yields are linear in $s_{1,t}$ and $s_{2,t}$. Details can be found in Pearson and Sun (1994) or Duffie (1996a).

Default is modeled as an unpredictable jump in a Poisson process. The intensity of the process for firm j at time t under the equivalent martingale measure is denoted $h_{j,t}$. In other words, the probability, under the equivalent martingale measure, that firm j defaults during the time $(t, t + dt)$, conditional on the firm not defaulting prior to t, is $h_{j,t}dt$. Throughout this article I use the terms "instantaneous probability of default" and "instantaneous default risk" interchangeably; they both refer to $h_{j,t}$. It should be emphasized that $h_{j,t}$ will not equal the true instantaneous probability of default as long as the market price of risk associated with the Poisson process is nonzero.

Now consider the price of a zero-coupon bond, issued by firm j, that pays a dollar at time T unless firm j defaults at or before T. In the event of default, the bond pays nothing. Denote the price of this bond as $V_j(t, T, 0, 0)$. (The third argument is the bond's coupon; the fourth is the recovery value in the event of default.) The price of this bond is given by the expectation, under the equivalent martingale measure, of an adjusted cumulative discount rate between t and T. This adjusted discount rate is $r_t + h_{j,t}$. Therefore the bond price is

$$V_j(t, T, 0, 0) = E_t^Q\left\{\exp\left[-\int_t^T (r_u + h_{j,u})du\right]\right\}. \tag{4}$$

The key modeling assumption in this article is the form of the default intensity process. I model the process $h_{j,t}$ as a translated single-factor square-root process plus two components tied to the default-free interest rate factors:

$$h_{j,t} = \alpha_j + h^*_{j,t} + \beta_{1,j}(s_{1,t} - \overline{s_{1,t}}) + \beta_{2,j}(s_{2,t} - \overline{s_{2,t}}); \tag{5}$$

$$dh^*_{j,t} = \kappa_j(\theta_j - h^*_{j,t})dt + \sigma_j\sqrt{h^*_{j,t}}dZ_{j,t}. \tag{6a}$$

Under the equivalent martingale measure, the process for $h^*_{j,t}$ can be written as

$$dh^*_{j,t} = (\kappa_j\theta_j - (\kappa_j + \lambda_j)h^*_{j,t})dt + \sigma_j\sqrt{h^*_{j,t}}d\hat{Z}_{j,t}. \qquad (6b)$$

This setup is designed to capture, in a tractable way, three important empirical features of corporate bond yield spreads. The most obvious is that the spreads are stochastic, fluctuating with the financial health of the firm. The $h^*_{j,t}$ process captures this fluctuation.

The second feature is that yield spreads for very high-quality firms are positive, even at the short end of the yield curve. This fact suggests that regardless of how healthy a firm may seem (i.e., how low $h^*_{j,t}$ is), there is some level below which yield spreads cannot fall, hence the α_j term is included in Equation (5). Nonzero yield spreads for very safe firms are likely a reflection of features unrelated to default risk, such as liquidity effects or state taxes. (Nonetheless, all these features are impounded into the "default risk" process $h_{j,t}$.) Such spreads can also reflect the risk that even the healthiest-looking firm can default if a sudden large event occurs, such as an unfavorable court decision exposing the firm to billions of dollars in claims.

The third feature is that yield spreads, especially spreads for lower quality bonds, appear to be systematically related to variations in the default-free term structure. A theoretical justification for this relation is in Longstaff and Schwartz (1995) and supporting empirical evidence is in Duffee (forthcoming). The stochastic process driving $h^*_{j,t}$ is assumed independent of the stochastic processes driving the default-free term structure; the correlation, if any, between default intensities and default-free interest rates is entirely captured by the $\beta_{1,t}$ and $\beta_{2,t}$ coefficients. This setup admits closed-form solutions for default-risky zero-coupon, no-recovery bond prices, in which the bonds' yields are linear in the variables $h^*_{j,t}$, $s_{1,t}$, and $s_{2,t}$. Details are contained in the Appendix.

A more comprehensive discussion of the link between default-risky and default-free term structures will be helpful in comparing this type of default-intensity model to a more structural model of default risk such as Longstaff and Schwartz (1995). Consider two hypothetical firms that have identical parameters in Equations (5), (6a), and (6b), except for their respective values of $\beta_{1,j}$. This difference will not affect the mean of $h_{j,t}$ because $\beta_{1,j}$ is multiplied by a mean-zero term.[3] The firms will then have identical mean values of $h_{j,t}$ equal to $\alpha_j + \theta_j$.

However, we could have written Equation (5) as follows:

$$h_{j,t} = \alpha^*_j + h^*_{j,t} + \beta_{1,j}s_{1,t} + \beta_{2,j}s_{2,t} \qquad (5')$$

[3] The term $\overline{s_{i,t}}$ is used in Equation (5) instead of θ_i (the long-run mean of $s_{i,t}$) for estimation-related reasons made clear in Section 3 of this article.

This alternative form is observationally equivalent to Equation (5) because the constant term α_j^* will absorb the $-\beta_{i,j}\overline{s_{i,t}}$ terms. But if we use Equation (5′) to perform the same comparative static exercise of varying $\beta_{1,j}$ while holding the other parameters fixed, we find that a higher value of $\beta_{1,t}$ corresponds to a higher mean $h_{j,t}$ as long as the mean of $s_{1,t}$ is positive.

The proper conclusion is that such a comparative static exercise is meaningless here. The model is not rich enough to allow us to infer from corporate bond data whether a firm with a high $\beta_{1,j}$ has a high mean $h_{j,t}$ because of this high $\beta_{1,t}$ [as in Equation (5′)] or because it happens to have a high α_j [as in Equation (5)].

By contrast, in Longstaff and Schwartz, there is an unambiguous relation between a firm's likelihood of default implied by bond prices and the correlation between default-free interest rates and the value of the firm's assets. In their model, a firm will have a lower implied likelihood of default if, holding all else equal, it has a more negative correlation between the value of its assets and the instantaneous interest rate. If this correlation is negative, an unexpected decline in the firm's assets (which pushes the firm closer to default) is typically accompanied by an increase in instantaneous interest rates, which pushes the growth rate of the firm's assets higher (thus pushing the firm further away from default). Therefore the value of the firm's assets is more stable if this correlation is negative than if this correlation is positive, and the likelihood that the firm defaults is correspondingly reduced.

The inability of the default-intensity model to capture the effect described by Longstaff and Schwartz is not the only limitation of this model. Equations (5) and (6a) allow for negative default intensities, either through a negative value of α_j or negative values of $\beta_{i,j}(s_{i,t} - \overline{s_{i,t}})$, $i = 1, 2$. The empirical section to follow documents that for the typical firm, α_j exceeds zero, but $\beta_{1,j}$ and $\beta_{2,j}$ are typically negative. Thus if default-free interest rates are sufficiently high, default intensities will be negative. This conceptual problem is a fairly common one in term structure modeling (both default-free and default-risky), and is largely ignored if the model accurately prices the relevant instruments.

Recall that Equation (4) is the price of a zero-coupon corporate bond that pays nothing in the event of default. However, modeling payments to bondholders in the event of default is of great practical importance because the amounts recovered are large. Moody's finds that, on average, senior unsecured bondholders receive approximately 44% of par in the event of default. Consider a senior unsecured zero-coupon bond issued by firm j that promises to pay, at time T, one dollar. This article follows Jarrow et al. (1997), among others, in assuming that the bond pays a fixed dollar value δ at time T in the event of default at or prior to T.

This assumption is equivalent to assuming that in the event of default, the bond pays, at the time of default, a fixed fraction δ of an otherwise equivalent default-free bond. For coupon bonds (the focus of this article), the price of such a default-free bond will, on average, be close to par. Therefore, in line with Moody's evidence, I assume that $\delta = 0.44$.

The price of a zero-coupon bond with recovery rate δ is denoted $V_j(t, T, 0, \delta)$. The absence of arbitrage requires that after default, the price of this bond equals the price of a default-free zero-coupon bond paying δ and maturing at time T. In addition, no arbitrage requires that, prior to default,

$$V_j(t, T, 0, \delta) = \delta P(t, T, 0) + (1 - \delta) V_j(t, T, 0, 0). \tag{7}$$

Therefore zero-coupon bonds with recovery value in the event of default can be easily priced in terms of default-free and default-risky, no-recovery zero-coupon bonds. The empirical work that follows uses bonds that pay coupons semiannually. The prices of such bonds are simply the sum of the prices of the individual cash flows associated with the bonds, although the yields on such bonds are no longer linear functions of the state variables. Details are in the Appendix.

Given that this article is looking at primarily investment-grade firms, the assumption of a constant recovery rate is harmless because it is an arbitrary identification assumption. If δ were allowed to vary stochastically over time and firms, the empirical tests that follow could not distinguish between variations in δ and variations in $h_{j,t}$ when the true $h_{j,t}$ is small.[4] This indeterminancy is explicitly recognized in Duffie and Singleton (1995, 1997), who model the product $(1-\delta)h$ instead of separately modeling δ and h. The inability to observe δ directly provides another reason (other than liquidity and tax-related differences) to be cautious in literally interpreting h as the equivalent martingale probability of instantaneous default.

2. Data Description

Corporate bond data are taken from the Lehman Brothers Fixed Income Database, which consists of month-end data on the bonds that make up the Lehman Brothers bond indexes. This dataset, which is discussed in more detail in Warga (1997) and Duffee (forthcoming), covers primarily investment-grade firms. I consider only noncallable, nonputable, senior unsecured straight bonds with semiannual coupons, no variation in promised coupon payments over time, no sinking fund provisions, and original maturities of under 35 years. All bonds in the dataset have at least one year remaining to maturity. Bond price observations that are calculated using a matrix are dropped. Therefore all prices are indicative bid prices quoted by

[4] For very risky instruments (large $h_{j,t}$), the assumption of $\delta = 0.44$ is restrictive because Equation (7) imposes a lower bound on corporate bond prices of $\delta P(t, T, 0)$. This bound can be violated for an instrument issued by a firm near default for which the recovery rate is actually substantially less than 0.44. I am grateful to a referee for pointing this out. Also note that Madan and Unal (1994) use prices of multiple classes of banks' obligations to separate δ from h.

Table 1
Summary statistics for 161 firms' noncallable bonds

	Across 161 firms		
Firm-level statistic	Minimum	Median	Maximum
Months of data	38	92	132
Mean number of fitted bonds per month	1.37	2.47	7.36
Mean years to maturity of fitted bonds	1.97	7.22	22.54
Minimum years to maturity of fitted bonds	1.00	1.04	14.29
Maximum years to maturity of fitted bonds	3.00	14.96	29.96
Mean coupon of fitted bonds	0	8.75	12.06

Each firm has at least 36 month-ends during January 1985 through December 1995 for which yields on at least two noncallable bonds were priced by Lehman Brothers traders.

traders. Relatively few noncallable bonds were issued prior to 1985, hence I examine only the period beginning January 1985 and ending December 1995.

I use a sample of 161 firms, where a firm is defined as any entity with a distinct six-digit CUSIP.[5] There are 96 industrial firms, 50 financial firms, and 15 utilities in this group. Each firm has at least 36 months in which at least two trader-quoted bond prices are observed. The list of firms is available from the author.

Firms sometimes issue a number of bonds with very similar maturity dates. Yields on such bonds are typically closely linked. To reduce the size of the estimation problem, in these cases I use a selection of bonds. To illustrate the selection procedure, consider a firm's bonds outstanding as of some month. Of the firm's bonds with a remaining maturity in the range [n years, $(n + 1)$ years), n an integer, all but the most recently issued bond is dropped. If multiple bonds share the same most recently issued date, all but the shortest-maturity bond is dropped. After dropping these similar-maturity bonds, each of the 161 firms considered here has at least 36 months in which at least two trader-quoted bond prices are observed, and many of the firms also have months during which a single trader-quoted bond price is observed.

Summary statistics about the bond data used in this article are presented in Table 1. The median number of observations for each firm is 92. Eleven of the firms have valid bond prices observed in every month, for a total

[5] Parents and their subsidiaries may have different six-digit CUSIPs. To take the most extreme example, my sample includes Ford Motors, Ford Motor Credit, Ford Capital B.V., and Ford Holdings as separate firms.

of 132 monthly observations. For any one firm, the number of bond prices observed varies over time. The mean number of bond prices (computed over those months for which the firm has at least one bond price) is calculated for each firm j, and summary statistics about this firm-specific mean are reported in the second row of Table 1. Across all firms, the lowest mean number of bonds is 1.4 per month, while the median is 2.5 and the maximum is 7.4. Therefore, although across all firms and all months this article uses 40,270 bond prices, the term structure of any one firm is estimated with a relatively small number of bonds. This paucity of data will be reflected in substantial uncertainty in the parameter estimates.

Table 1 also reports that for the median firm, the minimum bond maturity is 1 year, the mean is 7 years, and the maximum is 15 years. The median firm's mean coupon is 8.8%, paid semiannually. Two firms issued only zero-coupon bonds.

The Treasury yields are from the Center for Research in Security Prices (CRSP). At each month-end from January 1985 through December 1995, prices of seven different noncallable Treasury instruments are observed. They are the second most recently issued bills with maturities of 6 and 12 months, and the second most recently issued notes and bonds with maturities of 2, 3, 5, 10, and 30 years.[6] The most recently issued instruments are not included so as to avoid capturing any special liquidity premium that such instruments can have. In particular, these instruments (also known as on-the-runs) are often on special in the repo market, which distorts their observed yields. See Duffie (1996b) for a discussion.[7]

3. Estimation Methodology

There are a number of methods that can be used to estimate this model using both the cross-sectional and time-series properties of bond yields. I adopt the extended Kalman filter approach used by Chen and Scott (1995), Geyer and Pichler (1996), and Duan and Simonato (1997) in their analyses of Treasury yields.

[6] There are two slight complications. First, the computer programs used to conduct the estimation require that a bond have a constant coupon over time. Some Treasury bonds have odd first coupons. If the second most recently issued bond has an odd first coupon and has not passed the first payment date, the third most recently issued bond is used (with the same caveat for that bond). The second complication is that 30-year bonds appear in the sample beginning with February 1985, and there is not a second most recently issued 30-year bond until August. Until that time, 20-year bonds are used.

[7] Because on-the-runs trade more frequently than off-the-runs, it is possible that the use of second most recently issued bonds can introduce noise owing to stale prices. I looked for evidence of such noise by estimating my model of default-free yields twice; first with on-the-runs and then with second most recently issued bonds. I then compared the two sets of root mean squared measurement errors in bond yields. There was no evidence that the errors were smaller in magnitude for on-the-runs than for the bonds I examine in this article.

Firm j's bond prices depend on the processes for $s_{1,t}$ and $s_{2,t}$ as well as the process for $h^*_{j,t}$. Thus it would be more efficient to estimate the default-free interest rate and default probability processes jointly instead of separately. Unfortunately it is infeasible to jointly estimate the default-free processes and all 161 default probability processes. An alternative is for each firm to jointly estimate the default-free processes and the default-risky process for that firm. However, this alternative has the unsatisfying feature that it produces as many different estimated default-free processes as there are firms. I therefore estimate the default-free processes once, using only Treasury yields. I then use the resulting parameter estimates and corresponding predicted values of the unobserved state variables $s_{1,t}$ and $s_{2,t}$ to separately estimate the parameters of each firm's instantaneous default process. In other words, 162 different extended Kalman filter problems were solved, one for the default-free term structure and one for each firm's term structure.

3.1 Default-free parameter estimation

At time t we observe Treasury yields $Y_t = (y_{1,t}, \ldots, y_{7,t})'$. Denote the unobserved state variables by $S_t = (s_{1,t}, s_{2,t})'$. Suppressing the dependence of the model on the parameters to be estimated, the measurement and transition equations are

$$Y_t = z(S_t) + \epsilon_t, \quad E_{t-1}\left(\epsilon_t \epsilon_t'\right) = H$$

$$S_t = d + T S_{t-1} + \eta_t, \quad E_{t-1}\left(\eta_t \eta_t'\right) = R(S_{t-1}).$$

The function $z(S_t)$ maps the two state variables into seven yields. For zero-coupon instruments (the bill yields), this mapping is linear. For a coupon bond, this mapping is implicitly given by numerically solving for the yield corresponding to the coupon bond price implied by S_t. The H matrix is an 7×7 diagonal matrix with $H_{i,i} = \Sigma_i$. The transition equation components d and T are given by

$$d = \begin{pmatrix} \theta_1(1 - e^{-\kappa_1/12}) \\ \theta_2(1 - e^{-\kappa_2/12}) \end{pmatrix}, \quad T = \begin{pmatrix} e^{-\kappa_1/12} & 0 \\ 0 & e^{-\kappa_2/12} \end{pmatrix} \tag{8}$$

and $R(S_{t-1})$ is a 2×2 diagonal matrix with elements

$$R_{i,i}(S_{t-1}) = \kappa_i^{-1}\sigma_i^2[s_{i,t-1}(e^{-\kappa_i/12} - e^{-2\kappa_i/12}) + (\theta_i/2)(1 - e^{-\kappa_i/12})^2].$$

This notation introduces seven new variables Σ_i, $i = 1, 7$, which are the variances of the measurement errors of the yields.

The nonlinear function $z(S_t)$ is linearized around the 1-month-ahead forecast of S_t and $R(S_{t-1})$ is evaluated at the contemporaneous prediction of S_{t-1} (i.e., the recursion's prediction of S_{t-1} given bond yields through $t-1$).[8] The standard Kalman filter recursion is then used. Because $z(S_t)$ is nonlinear and the transition equation innovations η_t are nonnormal, the recursion's prediction of S_{t-1} conditioned on month $t-1$ information is not the conditional expectation $E(S_{t-1}|Y_{t-1}, Y_{t-2}, \ldots)$. Therefore the recursion's 1-month-ahead variance of S_t, which depends on the recursion's prediction of S_{t-1}, is not the correct 1-month-ahead variance of S_t. Thus the usual equivalence between the Kalman filter and quasi-maximum likelihood (QML) does not hold. Nonetheless, Monte Carlo evidence in Duan and Simonato (1997) suggests the procedure is reliable in the context of this class of models. Their article also contains a more detailed discussion of the properties of this methodology.

The Kalman filter requires an initial distribution of state variables. The unconditional distribution is used for the default-free process, which implies that the mean-reversion parameters κ_1 and κ_2 are required to be positive (so the unconditional distribution exists).

3.2 Default intensity parameter estimation

A firm's bond prices depend on both the firm's default intensity process and the default-free process. To estimate the default intensity process, the default-free process, estimated as described above, is assumed to be the true process. In addition, the values of $s_{i,t}$, $i = 1, 2$ predicted by the Kalman filter recursion (the smoothed estimates, based on information through the entire sample) are assumed to equal the true values of $s_{i,t}$, $i = 1, 2$. Denote these predicted values by $\hat{s}_{i,t}$. Then for any firm j, the only unobserved variable is $h^*_{j,t}$. An extended Kalman filter procedure is used to extract information about this variable.

Firm j has bond yields observed over a set of months in [January 1985, December 1995]. The means of $\hat{s}_{i,t}$, $i = 1, 2$ are computed over this set of months and are used as $\overline{s_{i,t}}$, $i = 1, 2$ in Equation (5). I use this sample mean instead of θ_i, $i = 1, 2$ so that the means of the last two terms in Equation (5) are identically zero for all firms j. This simplifies the interpretation of α_j.

Consider a given firm j. At time t, we observe $N_{j,t}$ corporate bond yields $Y_{j,t} = (y_{j,1,t}, \ldots, y_{j,N_{j,t},t})'$. The prior observation of the firm's bond prices occurred at $t - \tau$. Owing to missing observations, τ need not equal 1 month. The measurement and transition equations are (again suppressing

[8] It is possible for the predicted value of S_{t-1} to have a negative component, which is inconsistent with its stochastic process. If this occurs, the algorithm replaces this negative component with zero. This situation did not arise at the final parameter estimates, although it occasionally occurred at other sets of parameter estimates that were attempted during the optimization routine.

dependence on the parameters to be estimated)

$$Y_{j,t} = z_{j,t}(h^*_{j,t}; \hat{s}_{1,t}, \hat{s}_{2,t}) + \epsilon_{j,t}, \quad E_{t-\tau}\left(\epsilon_{j,t}\epsilon'_{j,t}\right) = H_{j,t}$$

$$h^*_{j,t} = d_j(\tau) + T_j(\tau)h^*_{j,t-\tau} + \eta_{j,t}, \quad E_{t-\tau}\left(\eta^2_{j,t}\right) = R_j(h^*_{j,t-\tau}, \tau).$$

The function $z_{j,t}(h^*_{j,t})$ maps $h^*_{j,t}$ into N_t yields. This mapping is implicitly given by numerically solving for the yield corresponding to the coupon bond price implied by $h^*_{j,t}$, $s_{i,t} = \hat{s}_{1,t}$, and $s_{2,t} = \hat{s}_{2,t}$. The $H_{j,t}$ matrix is an $N_t \times N_t$ diagonal matrix with each diagonal element equal to Σ_j, which is the common measurement error variance of the yields. A common error variance is used because for a given firm, both the number of bonds and their maturities vary over time. The transition equation components $d_j(\tau)$ and $T_j(\tau)$ are one-dimensional versions of Equation (8) with κ_i replaced with $\tau\kappa_j$. The conditional variance R_j is given by

$$R_j(h^*_{j,t-\tau}, \tau) = \kappa_j^{-1}\sigma_j^2[h^*_{j,t-\tau}(e^{-\tau\kappa_j/12} - e^{-2\tau\kappa_j/12}) + (\theta_j/2)(1 - e^{-\tau\kappa_j/12})^2].$$

Unlike the assumption made for the default-free process, I do not assume that each firm's default intensity process is stationary. Therefore I cannot use the unconditional distribution of $h^*_{j,t}$ to initiate the Kalman filter recursion. Instead, I use a least-squares approach to extract an initial distribution from the firm's first month of yield observations. Denote this first month as month 0. Then

$$z_{j,0}(h^*_{j,0}) \approx z_{j,0}(\theta_j) - Z_{j,0}\theta_j + Z_{j,0}h^*_{j,0},$$

where $Z_{j,0}$ is the linearization of $z_{j,0}$ around θ_j:

$$Z_{j,0} = \left.\frac{\partial z_{j,0}(h^*_{j,0})}{\partial h^*_{j,0}}\right|_{h^*_{j,0}=\theta_j}$$

Assuming that this linearization is exact, we can write the measurement equation for the firm's month 0 bond yields as

$$Y_{j,0} = z_0(\theta_j) - Z_0\theta_j + Z_{j,0}h^*_{j,0} + \epsilon_{j,0}.$$

This equation is rewritten in terms of $h^*_{j,0}$:

$$h^*_{j,0} = \frac{Z'_{j,0}(Y_{j,0} - z_0(\theta_j) + Z_0\theta_j)}{Z'_{j,0}Z_{j,0}} - \frac{Z'_{j,0}\epsilon_{j,0}}{Z'_{j,0}Z_{j,0}}.$$

Therefore the distribution of $h_{j,0}$ is assumed to have mean $Z'_{j,0}(Y_{j,0} - z_0(\theta_j) + Z_0\theta_j)/(Z'_{j,0}Z_{j,0})$ and variance $\Sigma_j/(Z'_{j,0}Z_{j,0})$. Given this initial distribution of the unobserved variable, the extended Kalman filter recursion proceeds as described for the default-free process.

Table 2
Extended Kalman filter estimates of a two-factor square-root model of Treasury bond yields, January 1985–December 1995

i	κ_i	θ_i	λ_i	σ_i^2	$\kappa_i\theta_i$	$\kappa_i + \lambda_i$
1	0.474	1.003	−0.011	0.000180	0.47491	0.463
	(0.032)	(0.048)	(0.004)	(0.000034)	(0.04184)	(0.032)
2	0.032	0.060	−0.015	0.002017	0.00192	0.017
	(0.009)	(0.056)	(0.021)	(0.001592)	(0.00160)	(0.016)

Bond maturity	Measurement error variance ($\times 10^6$)	Mean error (actual − fitted, basis points)	$\sqrt{\text{Mean square error}}$ (basis points)
6 months	12.189	−14.52	34.56
	(1.444)		
1 year	4.199	−12.80	20.11
	(0.423)		
2 years	0.101	0.46	1.80
	(0.049)		
3 years	0.229	0.31	4.20
	(0.051)		
5 years	0.720	1.75	8.08
	(0.157)		
10 years	0.297	−1.83	4.03
	(0.087)		
30 years	0.781	1.50	7.77
	(0.157)		

The instantaneous interest rate is

$$r_t = \alpha_r + s_{1,t} + s_{2,t}.$$

The dynamics are

$$ds_{i,t} = \kappa_i(\theta_i - y_{i,t})dt + \sigma_i\sqrt{s_{i,t}}dZ_{i,t} \qquad \text{(true measure)}$$
$$ds_{i,t} = (\kappa_i\theta_i - (\kappa_i + \lambda_i)y_{i,t})dt + \sigma_i\sqrt{s_{i,t}}d\hat{Z}_{i,t} \qquad \text{(martingale measure)}.$$

Month-end yields of the second most recently issued 6-month and 12-month Treasury bills and 2-, 3-, 5-, 10-, and 30-year coupon bonds are observed with normally distributed measurement errors independent across time and instruments. The value of α_r is fixed at −1.00 because, although the estimated value was below this value, the improvement in fit given by a smaller value was negligible. The standard errors (in parentheses) are computed assuming the Kalman filter linearization is exact.

4. Results

4.1 The default-free process

Estimation results for the default-free process are displayed in Table 2. Although the gradient of the Kalman filter likelihood makes it clear that α_r is negative, the data are unable to pin down the value of α_r with any reliability. Pearson and Sun (1994) also faced this problem using different data and a different estimation technique. I follow their approach and arbitrarily set α_r equal to a lower bound. I use $\alpha_r = -1$ because the improvement in fit given by substantially lower values of α_r was minimal.

Standard errors for the resulting restricted model are computed as with QML, although their validity is questionable given the nonlinearities. The first factor is closely related to (the negative of) the spread between long and short bond yields. The correlation between the predicted values of $s_{1,t}$ (based on information through the entire sample) and the spread between the 30-year bond yield and the 6-month Treasury bill yield is -0.97, while the correlation between first differences of these two series is -0.84. This factor exhibits significant mean reversion; its half-life is slightly less than 1.5 years. The second factor is closely related to long bond yields. The predicted values of $s_{2,t}$ (again, based on information through the entire sample) move almost in lockstep with the Treasury's 30-year bond yield. (The correlation of first differences is 0.97.) With a half-life of over 20 years, this factor is close to a martingale.

The interpretation of one factor as the slope of the Treasury yield curve and the other factor as the level of Treasury yields is standard. Chen and Scott (1995), Geyer and Pichler (1996), and Duan and Simonato (1997) produce similar results using the same technique over different periods, while Chen and Scott (1993) produce similar results using a different technique over a different sample period.

Because a negative value of α_r implies the possibility of negative interest rates, it is worth exploring what features of the data generate this negative estimate. If α_r is zero (or positive), the model is limited in its ability to generate both steeply sloped term structures and flat, low term structures with reasonable volatility behavior.

We explain this limitation in two steps. First, consider how the slope of the term structure is related to the state variable associated with (the negative of) the term structure, $s_{1,t}$. When the term structure is flat, $s_{1,t}$ is larger than it is when the term structure is steeply sloped. Recall that $s_{1,t}$ cannot be negative, regardless of the slope of the term structure. Therefore, if the model fits accurately the slope of the term structure both when it is steeply sloped and when it is flat, $s_{1,t}$ will be relatively large at those times when the term structure is flat.

Now consider how well the model will fit yields when the term structure is both flat and low. If $\alpha_r = 0$, the model can get the slope right, but may be unable to match the level. The observed yields may be too low to be consistent with the large value of $s_{1,t}$ necessary to fit the slope, even if $s_{2,t}$ is set to zero. Duffie and Singleton (1997) encounter this problem. Even if there is a value of $s_{2,t}$ consistent with the low yields, this value might be very small. In this square-root diffusion model, a small value of $s_{2,t}$ implies long maturity yields have very low volatility, which is a strong restriction on the model.

If a negative constant term is included in the default-free process, $s_{2,t}$ can be large even when yields are low, allowing for a better fit. This flexibility is especially important over the period examined in this article, which exhibited both very steep term structures (in the early 1990s) and flat, very low term structures (in 1995). In the context of the model, as long as the constant term is sufficiently negative, the precise value of α_r is not particularly important. Therefore the estimation routine has great difficulty pinning it down.

This two-factor model does a rather poor job of pricing bills, but does reasonably well at pricing long-maturity instruments. As Table 2 indicates, fitted yield curves are consistently higher than actual yields at the short end. The mean error (actual less fitted) for 6- and 12-month bills are -14 and -13 basis points (b.p.), respectively. The root mean squared errors (RMSE) in the bill yields are between 20 and 34 b.p. For longer maturity instruments, mean errors are all less than 2 b.p. and RMSE's are all less than 9 b.p.

For the purposes of this article, the inability of the model to accurately price short-maturity instruments is not an important concern. Our ultimate goal is to price corporate bonds, and all of the corporate bonds examined in this article have more than 1 year to maturity. Thus over the range of maturities relevant for pricing corporate bonds, the default-free process estimated here adequately prices Treasury bonds.

Careful consideration of the estimated default-free rate process is beyond the scope of this article. However, there is evidence of misspecification beyond the failure of the model to fit accurately the short end of the yield curve. The model assumes that the process for $s_{1,t}$ is independent of the process for $s_{2,t}$, but the fitted values are negatively correlated. The correlation of monthly changes in these fitted values is -0.37.

4.2 The default intensity processes

Table 3 summarizes the parameter estimates of the firms' default intensity processes. The table reports the median and interquartile range of each estimated parameter (there are 161 estimates of each parameter, one for each firm). It also reports the median and interquartile range of three firm-specific values: firm j's mean $\hat{h}_{j,t}$, the firm's mean $\hat{h}^*_{j,t}$, and the square root of the mean of the squared differences between the actual and fitted yields on firm j's bonds (RMSE). In contrast to the way the fitted state vector was defined for the default-free process, here the fitted value of $h_{j,t}$ is the filtered prediction from the Kalman filter (the prediction based on information through t), not the smoothed prediction (information through the entire sample). This choice is made to simplify the specification tests performed in Section 6.

Before discussing the individual parameter estimates in detail, it is worth noting that the differences between the parameters' lower-quartile estimates and higher-quartile estimates are quite large. Although it is possible that these large differences reflect substantial cross-firm variation in the true parameters, another likely explanation is imprecision in the estimates. As noted earlier, any one firm's term structure is fitted using relatively few

Table 3
Summary of extended Kalman filter estimates of 161 firms' default risk processes implied by corporate bond yields

Variable	1st quartile	Median	3rd quartile
$100 \cdot \alpha_j$	0.396	0.749	1.175
κ_j	0.023	0.238	0.600
$100 \cdot \theta_j$	0.072	0.559	2.814
λ_j	−0.485	−0.235	−0.050
σ_j	0.051	0.074	0.104
$\beta_{1,j}$	−0.242	−0.096	0.000
$\beta_{2,j}$	−0.080	−0.009	0.062
$100 \cdot \kappa_j\theta_j$	0.020	0.103	0.306
$\kappa_j + \lambda_j$	−0.150	−0.033	0.118
$100 \cdot \sqrt{\Sigma_j}$	0.082	0.101	0.121
Mean fitted $h_{j,t} \cdot 100$	1.086	1.359	1.876
Mean fitted $h^*_{j,t} \cdot 100$	0.238	0.573	1.079
RMSE of yield to maturity (basis points)[a]	7.39	9.83	11.05

The instantaneous default-free interest rate is given by $r_t = \alpha_r + s_{1,t} + s_{2,t}$, where $s_{1,t}$ and $s_{2,t}$ are independent square-root processes. Firm j's instantaneous default risk is given by

$$h_{j,t} = \alpha_j + h^*_{j,t} + \beta_{1,j}(s_{1,t} - \overline{s_{1,t}}) + \beta_{2,j}(s_{2,t} - \overline{s_{2,t}}),$$

where $h^*_{j,t}$ follows a square-root process that is independent of the processes for $s_{i,t}$, $i = 1, 2$:

$$dh^*_{j,t} = \kappa_j(\theta_j - h^*_{j,t})dt + \sigma_j\sqrt{h^*_{j,t}}dZ_{j,t} \qquad \text{(true measure)}$$

$$dh^*_{j,t} = (\kappa_j\theta_j - (\kappa_j + \lambda_j)h^*_{j,t})dt + \sigma_j\sqrt{h^*_{j,t}}d\hat{Z}_{j,t} \qquad \text{(martingale measure)}.$$

Month-end yields on firm j's noncallable coupon bonds are all observed with i.i.d., normally distributed measurement error with mean zero and variance Σ_j. The estimation period is January 1985–December 1995, although most of the firms do not have data over this entire period.
[a] The square root of the mean of the squared differences between the actual and fitted yields to maturity on firm j's bonds.

bonds, and no firm has more than 11 years of monthly observations. Therefore the individual parameter estimates are subject to substantial uncertainty. Accordingly, in this discussion I focus on the median parameter estimates across all 161 firms.

The median estimate of the constant term α_j is 0.0075. To get a sense of how this term affects yield spreads, assume that the default-free states $s_{i,t}$, $i = 1, 2$ are at their mean fitted values. Then, regardless of how healthy the firm is (at the extreme, $h^*_{j,t} = 0$), this estimate implies that the yield spread on a near-zero-maturity zero-coupon bond can be no lower than 41.9 b.p., points, assuming a 44% recovery rate.

As discussed in Section 1, it is tempting to interpret this term as a constant bond premium unrelated to default risk. But it is important to note that the positive estimates of α_j are not based on any liquidity premium or state tax effect imposed on the model. The features of the data that give rise to these positive estimates will be discussed in the next section.

The median estimate of κ_j is 0.24. Thus, even though stationarity is not imposed, default intensities appear to be stationary, with a half-life of less than 3 years.[9] The median estimate of λ_j is negative, implying that investors require compensation for variations in default risk. Put another way, assets that pay off when firms' default probabilities rise are highly valued by investors. The estimates of λ_j are typically larger in absolute value than those of κ_j, hence the median estimate of mean reversion parameter under the equivalent martingale measure, $\kappa_j + \lambda_j$, is negative. Negative mean reversion implies that under the equivalent martingale measure, instantaneous default risk is explosive, in the sense that higher levels of risk correspond to more rapid growth in risk. The next section explains what features of the data generate these negative estimates.

Table 3 also documents a modest negative relation between default intensities and the default-free term structure. The median estimate of $\beta_{1,j}$ implies that a 100-b.p. decline in $s_{1,t}$ lowers $h_{j,t}$ by 0.00096, which, given a recovery rate of 44%, corresponds to a decline in yield spreads on near-zero-maturity instruments of 5.4 b.p. The median estimate of $\beta_{2,j}$ is economically close to zero. The estimate implies that a 100-b.p. decline in $s_{2,t}$ corresponds to a decline in yield spreads on near-zero-maturity instruments of only 0.5 b.p.

Also shown in Table 3 is the median firm's mean fitted instantaneous default probability, which is 1.4% per annum. The constant term accounts for most of this default risk; the median firm's mean value of $\hat{h}^*_{j,t}$ is less than 0.6% per annum. It is worth noting that default probabilities based on observations of actual defaults are smaller than even the latter value. The bulk of the firms in my sample are rated either A or Baa by Moody's. Fons et al. (1994) report the historical probability of a Baa-rated firm defaulting within 5 years is only 2%, while the historical probability for an A-rated firm is less than 1%.

The final point to take from Table 3 is that this model fits corporate bond prices well. The median estimate of the standard deviation of measurement error in bond yields, $\sqrt{\Sigma_j}$, is only 10 b.p. Similarly, the median root mean squared error (RMSE) of yields is slightly over 9 b.p., and for three-fourths of the firms the RMSE is no greater than 11 b.p.

5. Results by Credit Rating

In this section we break down the model's estimates according to the credit rating of the bonds' issuers. There are two reasons for this. First, an examination of the disaggregated results will help evaluate the consistency of the

[9] The estimates of κ_j can be used to illustrate the great uncertainty in a single firm's parameter estimates. Standard errors for each firm's estimate of κ_j were constructed using the standard QML assumptions and assuming that the estimates from the default-free process (both the parameters and the values of $s_{i,t}$, $i = 1, 2$) were known. The median standard error was 0.309.

model in a way that cannot be captured through analyses of bond pricing errors. Second, this examination will help explain why the data produce estimates of α_j that are typically positive and estimates of $\kappa_j + \lambda_j$ that are typically negative. We shall see that the signs of these parameter estimates are crucial to fitting basic features of the data.

To understand how this disaggregation conveys information about the consistency of the model, consider some firm j with a credit rating of, say, Aa at time t. (Following standard practice, we define a firm's credit rating as the rating on the firm's senior unsecured bonds.) Imagine that over time the firm's credit quality declines such that at time $\tau > t$ the firm's credit rating is Baa. The model of default risk used here says that this change in credit quality is impounded entirely in the state variables $h^*_{j,t}$, $s_{1,t}$, and $s_{2,t}$. The parameters of the default intensity process in Equations (5), (6a), and (6b) are fixed as the firm's credit quality changes.

We might think to test directly for instability in this firm's default intensity parameters. However, for any single firm, the parameters are estimated with so little precision that tests of parameter instability will have minimal power. An indirect test is to look for systematic variation in estimated parameters across firms with different credit ratings. If there were any systematic relation, it would suggest that as a firm's credit quality changed, the parameters of its default intensity process would also change.

This argument is not foolproof. The bond pricing model used here is consistent with a systematic relation between parameters and credit ratings as long as this relation is determined prior to the first time bond prices are observed. For example, firms with default intensity parameters that fall into a certain range may typically choose to set their capital structure so that their initial default intensity $h^*_{j,0}$ is relatively high. But if the firm alters its choice when it has outstanding bonds, the model's parameters will be unstable.

The results displayed in Table 4 tell us what a firm's credit rating actually conveys about the firm's default intensity process. Credit ratings are from Moody's unless a bond had no Moody's credit rating in the dataset. In this case an S&P rating was used. Panel A of the table reports median estimates of the parameters of firms' default probability processes, sorted by each firm's "initial" credit rating, which is defined as a firm's credit rating as of the first month of data used to estimate the instantaneous default process for the firm.

A graphical representation of Table 4, panel A is in Figure 1. This figure displays the term structure of yield spreads for three hypothetical firms, where the default intensity parameters of the firms are given by the median estimates in panel A for Aa-, A-, and Baa-rated firms. A term structure for Aaa-rated firms is not included because their median estimate of θ_j is improbably high. (The implied Aaa spread curve starts below the other spread curves, but quickly rises above the curves for Aa- and A-rated firms.)

Table 4
The relation between credit ratings and estimated default probabilities

Panel A: Median parameter estimates grouped by firms' initial credit ratings

Variable	Aaa	Aa	A	Baa
Number of firms	7	39	79	36
$100 \cdot \alpha_j$	0.594	0.637	0.739	0.961
κ_j	0.269	0.186	0.242	0.212
$100 \cdot \theta_j$	1.730	0.499	0.559	0.628
λ_j	−0.130	−0.216	−0.257	−0.307
σ_j	0.077	0.074	0.078	0.059
$\beta_{1,j}$	−0.056	−0.077	−0.090	−0.171
$\beta_{2,j}$	0.000	0.001	−0.033	−0.006
Mean fitted $h^*_{j,t} \cdot 100$	0.337	0.440	0.537	0.864
RMSE of yield to maturity (basis points)	6.95	8.41	9.38	9.76

Panel B: Estimated instantaneous default risk behavior, sorted by contemporaneous credit rating

Month t credit rating	Number of observations	Median values (×100) of: $\hat{h}_{j,t}$	$\sigma_j(\hat{h}^*_{j,t})^{1/2}$	$[\widehat{Var}_{t-1}(\hat{h}^*_{j,t})]^{1/2}$
Aaa	440	0.793	0.417	0.123
Aa	2,351	0.969	0.462	0.131
A	7,937	1.231	0.428	0.125
Baa	3,615	1.834	0.709	0.204
Ba	34	1.293	0.896	0.335

Parameters are defined in Table 3. In panel A, a firm's initial credit rating is defined as the credit rating on the firm's senior unsecured bonds in the first monthly observation used to estimate the firm's default process. In panel B, fitted monthly values are denoted with hats. The final column in panel B reports the 1-month-ahead variance of the fitted value implied by the Kalman filter recursion.

The instantaneous default risk $h^*_{j,t}$ is set to the appropriate median fitted $\overline{h^*_{j,t}}$ reported in the same table. The parameters of the default-free process are taken from Table 2 and the factors $s_{i,t}$, $i = 1, 2$ are set to their mean values over 1985 through 1995.

The spread curves in Figure 1 are consistent with stylized facts about term structures of yield spreads. Earlier research has shown that on average across investment-grade bonds, the term structure of yield spreads is upward sloped.[10] Litterman and Iben (1991), Fons (1994), and Duffee (forthcoming) find that this positive slope is stronger for lower-rated bonds than higher-rated bonds, which is also consistent with Figure 1. For example, Fons (1994) finds that yield spreads for Aaa-rated coupon bonds are essentially flat across maturities, while yield spreads for Baa-rated coupon bonds increase about 2 b.p. for each additional year to maturity.

[10] The focus here is on investment-grade bonds because the dataset has little information on speculative-grade bonds. The slopes of term structures of yield spreads for speculative-grade bonds are examined by Sarig and Warga (1989), Fons (1994), and Helwege and Turner (1997). On balance, the evidence concerning the slope of the typical speculative-grade term structure is mixed.

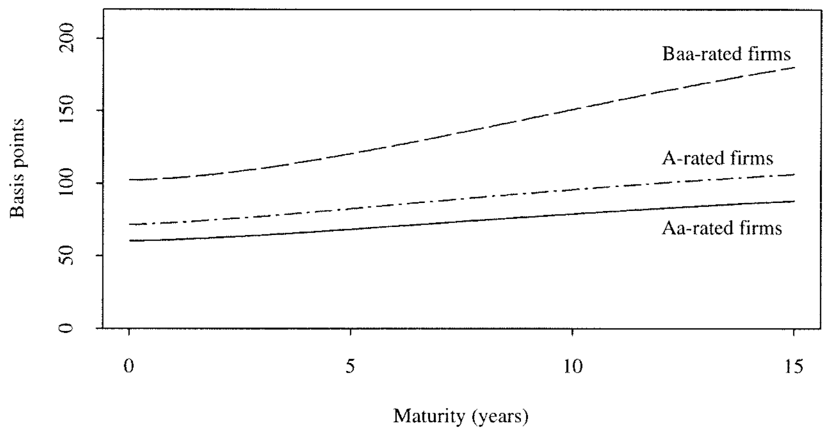

Figure 1
Term structures of yield spreads
This figure displays hypothetical term structures of zero-coupon corporate bond yields less zero-coupon Treasury bond yields implied by a two-factor translated square-root model of Treasury yields and a translated square-root model of instantaneous default risk. The default-risk parameters for each credit rating are the median estimates across those firms with the given rating. The two default-free factors are set equal to their mean values over the sample period 1985–1995.

In Table 4, the median estimate of a firm's mean fitted $h^*_{j,t}$ rises monotonically as the firm's credit rating falls. This is part, but not all, of the reason why lower-rated firms have higher yield spreads in Figure 1. Median estimates of α_j are also higher for lower-rated firms. For example, in Figure 1, the short end of the Aa-rated yield curve is 42 b.p. lower than the short end of the median Baa-rated yield curve. Of this amount, slightly more than half is the result of the difference in the two values of $h^*_{j,t}$ used to construct the curves. The remainder owes to the differences in α_j.

Similarly, part of the differences among the slopes of the various spread curves are explained by differences in $h^*_{j,t}$, with the remainder explained by variations in parameter estimates. In order for higher values of $h^*_{j,t}$ to correspond to steeper average yield spread slopes, $\kappa_j + \lambda_j$ must be negative. When this sum is negative, investors price corporate bonds as if they expect $h^*_{j,t}$ to rise over time, inducing a positive slope to the term structure of yield spreads.[11] Moreover, for a fixed negative $\kappa_j + \lambda_j$, a higher value of $h^*_{j,t}$ implies a steeper yield spread curve.

Although negative values of $\kappa_j + \lambda_j$ are important in capturing the variations of yield spread slopes as credit quality changes, they induce fairly steep yield curves for even high-quality firms. Roughly speaking, the role of a positive α_j is to dampen the overall steepness of the yield spread curves.

[11] This is strictly true only for a stochastic $h^*_{j,t}$. If $h^*_{j,t} = \kappa_j\theta_j = 0$, $h^*_{j,t}$ is fixed at zero and thus the yield spread will depend on only α_j and the two terms related to the default-free factors.

Without the α_j parameter, the model would generate yield spread curves for highly rated firms much steeper than those found in the data.

To see this clearly, imagine setting α_j to zero. Then $h^*_{j,t}$ and θ_j would have to be increased by α_j in order to hold both $h_{j,t}$ and expectations of future values of $h_{j,t}$ fixed. But under the equivalent martingale measure, this increase in $h^*_{j,t}$ would imply a more explosive default risk process, and thus a steeper yield spread curve. For example, consider the yield spread curve for Aa-rated firms in Figure 1, which rises from 60 b.p. at the short end to 88 b.p. at 15 years. If the estimate of α_j were set to zero while adjusting $h^*_{j,t}$ and θ_j by the appropriate amount, the resulting yield spread curve would rise from 60 b.p. to 134 b.p. over the same maturity span. Without $\alpha_j > 0$ and $\kappa_j + \lambda_j < 0$, this model could not simultaneously match two fundamental features of the data: low, nearly flat term structures when firms are very creditworthy and high, steeper term structures when firms are less creditworthy.[12]

Variations in estimates of $\kappa_j + \lambda_j$ across credit ratings also contribute to the pattern of steeper yield spread slopes for lower-rated firms. Because $\kappa_j + \lambda_j$ is more negative for the lower-rated firms, default risk is more explosive under the equivalent martingale measure for low-rated firms than for high-rated firms. For example, in Figure 1, the spread between the Aa-rated spread curve and the Baa-rated spread curve widens from 42 b.p. at the short end to 92 b.p. at 15 years. If the Aa-rated spread curve were constructed using the Baa-rated firms' median value of $\kappa_j + \lambda_j$, the spread between the curves at 15 years would be only 60 b.p. Thus most of the difference between the slopes of the Aa-rated spread curve and the Baa-rated spread curve owes to the difference in $\kappa_j + \lambda_j$.

The final point to take from panel A of Table 4 is lower-rated firms' greater sensitivity to the slope of the default-free term structure. The median estimates imply that if $s_{1,t}$ increases by 100 b.p., yield spreads on Aaa-rated, near-zero-maturity instruments fall by instantaneous default intensities by only 3 b.p., but yield spreads on otherwise equivalent Baa-rated instruments fall by nearly 10 b.p. Across all credit ratings, default intensities respond very weakly to changes in $s_{2,t}$. The general pattern of greater sensitivity for lower-rated firms is consistent with the empirical behavior of yield spreads documented in Duffee (forthcoming), but the magnitudes of the implied inverse relations between yield spreads and default-free interest rates reported here are uniformly weaker than reported there.

Panel A of Table 4 disaggregated parameter estimates by firms' initial credit ratings. A different type of disaggregation is used in panel B. Fitted values of $h_{j,t}$ are sorted according to firm j's time t credit rating instead of firm j's time 0 credit rating. As expected, these values are inversely related to

[12] This failure was confirmed in an earlier version of this article which did not include α_j.

firms' credit ratings. Firms rated Aaa have a median estimated instantaneous default probability of 0.79%; this median estimate rises to 1.83% for the Baa firms. (A few observations of Ba-rated firms are available, for which the median estimate of $h_{j,t}$ is less than that of Baa-rated firms. This anomaly is likely created by the sample size.)

The panel also reports a general pattern of higher default risk volatility for lower-rated firms. Volatility is measured in two ways: the implied instantaneous standard deviation of $h^*_{j,t}$, which is $\sigma_j\sqrt{h^*_{j,t}}$, and the 1-month-ahead standard deviations of the fitted values of $h_{j,t}$ produced by the Kalman filter recursion. Median values of both measures are lowest for Aaa-rated firms and highest for Baa-rated firms. (They are even higher for Ba-rated firms, but the limited number of observations makes these estimates unreliable.) This pattern is, however, violated by the relative ordering of the volatilities of Aa-rated and A-rated firms.

On balance, the disaggregated estimates from the default-risk model suggest that the model is misspecified in certain ways. Although part of the variation, across credit ratings, in term structures of yield spreads is simply the result of variations in $h^*_{j,t}$, the estimates also capture part of this variation in systematic variations in parameters across credit ratings. Of course, it is possible that the model is not misspecified. Instead, the median parameter estimates in Table 4 may be poor estimates of the true parameters for typical firms. More formal specification tests are considered in the next section.

6. Specification Tests

In this section we consider whether the stochastic process for $h^*_{j,t}$ specified in Equation (6a) is an accurate description of the behavior of $h^*_{j,t}$ implied by corporate bond prices. Recall that $\hat{h}^*_{j,t}$ is the Kalman filter's estimate of $h^*_{j,t}$ based on information through t. The filter also produces a one-step-ahead expectation and variance of this estimate. (One step is usually 1 month, but if there are missing values, a step will be multiple months.) Therefore we can construct normalized one-step-ahead innovations in $\hat{h}^*_{j,t}$, denoted $\zeta_{j,t}$. If the model of instantaneous default risk is correctly specified, neither $\zeta_{j,t}$ nor $|\zeta_{j,t}|$ should be forecastable using information dated prior to t.

6.1 Persistence

We first test for the presence of autoregressive components in $\zeta_{j,t}$ by fitting it to an AR(3). Note that under the null hypothesis, the estimated coefficients are equal (to zero) across all firms, as are the variances of the regression residuals. Therefore I can estimate the AR(3) equations jointly across all j:

$$\begin{aligned} \zeta_{j,t} &= b_0 + \sum_{i=1}^{3} b_i \zeta_{j,t-i} + e_{j,t}, \\ E(e_{j,t}) &= 0,\ E(e_{j,t}^2) = \Omega, \\ E(e_{i,t}e_{j,t}) &= \rho\Omega,\ i \neq j,\ E(e_{i,t}e_{j,t-m}) = 0, m \neq 0,\ \forall i, j. \end{aligned} \tag{9}$$

Under the null hypothesis that the model of instantaneous default risk is correctly specified, the AR coefficients should equal zero. I estimate Equation (9) simultaneously across all firms using QML. I assume that the residuals from the regression are contemporaneously correlated across firms and uncorrelated across time. The QML procedure maximizes the likelihood function associated with Equation (9), assuming that the errors are jointly normally distributed. The standard errors are then adjusted to account for nonnormality of the errors. In Equation (9), note that the coefficients are not indexed by j; the estimated parameters are assumed to be equal across all j.

The results of estimation of Equation (9) are reported in panel A of Table 5. Some specification error is evident. The first AR term is 0.22 and is statistically significantly different from zero. Therefore $\zeta_{j,t}$ exhibits persistence, indicating that estimated values of κ_j do not capture completely the mean reversion in $\hat{h}^*_{j,t}$. Of independent interest is the cross-correlation of innovations in firms' instantaneous default probabilities. The point estimate of ρ is 0.29, and as the t-statistic indicates, it is estimated very precisely. This point estimate implies that variations in firms' default probabilities have common components, but these common components are not very large compared to idiosyncratic components. This result is not altered by sorting firms by their senior unsecured credit ratings. In results that are not detailed here, I find that QML estimation of Equation (9) broken down by major credit rating (Aaa, Aa, etc.) produces similar estimates of ρ.

Panel A of Table 5 also reports a test of the accuracy of the volatility specification of Equation (6a). In Equation (6a), the only source of persistence in the volatility of $h^*_{j,t}$ is the persistence in the level of $h^*_{j,t}$. There are two issues of interest. First, are there persistent variations in volatility (GARCH-like effects) that are not accounted for by variations in the level of $h^*_{j,t}$? Second, does the square-root specification capture the true relation between $h^*_{j,t}$ and innovations in $h^*_{j,t}$? (Is the "level effect" correctly modeled?) I estimate Equation (10) jointly across all firms with QML:

$$|\zeta_{j,t}| = b_0 + \sum_{i=1}^{3} b_i |\zeta_{j,t-i}| + b_4 \hat{h}^*_{j,t-1} + e_{j,t}. \tag{10}$$

The volatility of $\zeta_{j,t}$ is measured with absolute values. This choice is based on Davidian and Carroll's (1987) conclusion that when distributions

Table 5
Specification tests of the default-risk model

Panel A: Forecasting default-risk innovations and their volatilities

Dependent variable	Explanatory variables: Lag 1	Lag 2	Lag 3	$\hat{h}^*_{j,t-1}$	ρ
$\zeta_{j,t}$	0.223	−0.008	−0.020	—	0.286
	(9.97)	(−0.58)	(−1.37)		(7.78)
$\lvert\zeta_{j,t}\rvert$	0.134	0.026	0.020	0.136	0.157
	(8.04)	(1.90)	(1.88)	(0.56)	(3.81)

Panel B: Relations between default risk and default-free interest rates

Month $t-1$ rating	Number of observations	Explanatory variables: $\Delta\hat{s}_{1,t}$	$\Delta\hat{s}_{2,t}$	$\Delta\hat{s}_{1,t-1}$	$\Delta\hat{s}_{2,t-1}$	ρ
Aaa	430	−11.19	3.03	9.73	13.61	0.208
		(−0.53)	(0.16)	(0.55)	(0.76)	(2.94)
Aa	2,308	−15.18	−17.81	16.50	27.84	0.302
		(−0.95)	(−1.04)	(1.13)	(1.74)	(8.36)
A	7,807	−7.18	−29.79	10.19	16.27	0.316
		(−0.46)	(−1.88)	(0.71)	(1.22)	(7.78)
Baa	3,536	−17.09	−23.87	2.44	13.40	0.295
		(−1.00)	(−1.27)	(0.17)	(0.82)	(7.40)

The stochastic component of firm j's instantaneous default risk that is independent of the default-free term structure is $h^*_{j,t}$. The fitted monthly value $\zeta_{j,t}$ is the normalized innovation in the fitted values of $h^*_{j,t}$, where the normalization uses the one-period-ahead variance in fitted $h^*_{j,t}$ implied by the Kalman filter recursion. In panel A, both $\zeta_{j,t}$ and $|\zeta_{j,t}|$ are fit to an AR(3). The fitted level $\hat{h}^*_{j,t-1}$ is included in the latter regression. There are 13,583 observations. In panel B, $\zeta_{j,t}$ is regressed on the current and lagged first differences of the fitted default-free interest rate factors. These differences are denoted $\Delta\hat{s}_{i,t}, i = 1, 2$. Observations are sorted by firm j's month $t-1$ credit rating and the regression is estimated separately for each group. The regressions are estimated jointly across firms with QML, where the correlation between different firms' regression errors is ρ, assumed to be constant across firms and time. Noncontemporaneous innovations are assumed independent. T-statistics are in parentheses.

are characterized by fat tails, it is often more efficient to estimate volatility functions using absolute values of residuals instead of squares. The variance-covariance matrix of the residuals from Equation (10) is assumed to have the same form as that for Equation (9).

The results in Table 5 indicate moderate GARCH-like effects in $\zeta_{j,t}$. The AR(1) coefficient for $|\zeta_{j,t}|$ is 0.13 and is strongly statistically significant. These results indicate that in the data there is some source of persistence in the volatility of $h^*_{j,t}$ that is not captured by the level of $h^*_{j,t}$. On the other hand, the results in Table 5 indicate that the square-root form of the level effect is an adequate characterization; the level of $h^*_{j,t-1}$ is insignificantly related to $|\zeta_{j,t}|$.

6.2 The relation between default risk and default-free rates

The model of default risk used in this article is designed to capture any correlation between default risk and default-free interest rates solely through the coefficients $\beta_{1,j}$ and $\beta_{2,j}$. The bond pricing formulas used to compute zero-coupon corporate bond prices rely on the assumption that $h^*_{j,t}$ is independent of the components of the default-free interest rate, $s_{1,t}$ and $s_{2,t}$. Here I investigate the appropriateness of that assumption.

I regress $\zeta_{j,t}$ on the contemporaneous and lagged changes in the fitted values of $s_{1,t}$ and $s_{2,t}$:

$$\zeta_{j,t} = b_0 + b_1\hat{s}_{1,t} + b_2\hat{s}_{2,t} + b_3\hat{s}_{1,t-1} + b_4\hat{s}_{2,t-1} + e_{j,t}. \qquad (11)$$

Observations are sorted by firm j's month $t-1$ credit rating and Equation (11) is estimated with QML separately for each group. The variance-covariance matrix of the residuals is again assumed to have the same structure as Equation (9).

The results are reported in panel B of Table 5. From the perspective of statistical significance, they are encouraging. Across all four credit rating groups, no coefficient is individually significant at the 5% level. In addition, the economic significance of the coefficients is small. For concreteness, consider the results for Baa-rated firms. Imagine that $s_{1,t}$ increases by 100 b.p. The point estimates in Table 5, combined with the median standard deviation for Baa-rated firms reported in panel B of Table 4, imply an increase in near-zero-maturity yield spreads of 2 b.p. The implied responsiveness of Baa yield spreads to changes in $s_{2,t}$ is greater, but the point estimates indicate that this responsiveness is quite transitory—most is reversed within a month.

6.3 Stale bond prices

Information that affects a firm's value should be immediately impounded into both the firm's stock price (if it has publicly traded stock) and its bond prices. But because the bond prices used in this article are not transaction prices, but rather indicative prices supplied by traders, it is possible that the bond prices respond to the information with a lag. If so, stock returns should predict future bond returns. Using a different dataset of indicative bond prices, Kwan (1996) documented that weekly changes in bond yields lagged stock returns. Here I examine the importance of lagged responsiveness at the monthly frequency.

To test for stale bond prices, I regress the monthly normalized innovation $\zeta_{j,t}$ on lags 0–3 of the firm's monthly stock return $\pi_{j,t}$. The stock returns are from CRSP and are available for 128 of the 161 firms:

$$\zeta_{j,t} = b_0 + b_1\pi_{j,t} + b_2\pi_{j,t-1} + b_3\pi_{j,t-2} + b_4\pi_{j,t-3} + e_{j,t}. \qquad (12)$$

Observations are sorted by firm j's month $t-1$ credit rating and Equa-

Table 6
The relation between firms' estimated instantaneous default probabilities and their stock returns

Credit rating	Number of observations	Stock return Lag 0	Lag 1	Lag 2	Lag 3	ρ
Aaa	124	−3.521	−0.876	−3.767	0.426	−0.275
		(−2.20)	(−1.06)	(−3.73)	(0.32)	(−2.06)
Aa	1,591	−0.859	−0.979	−0.685	−0.673	0.263
		(−2.12)	(−2.29)	(−2.01)	(−2.04)	(7.48)
A	5,781	−0.640	−0.655	−0.577	−0.338	0.299
		(−2.40)	(−3.25)	(−3.34)	(−2.06)	(8.34)
Baa	2,993	−1.424	−0.877	−0.767	−0.613	0.247
		(−4.18)	(−3.22)	(−2.59)	(−2.81)	(6.96)

The stochastic component of firm j's instantaneous default risk that is independent of the default-free term structure is $h^*_{j,t}$. Normalized innovations in the fitted values of $h^*_{j,t}$ are regressed on current and three lags of firm j's monthly stock returns, where the normalization uses the 1-month-ahead variance in fitted $h^*_{j,t}$ implied by the Kalman filter recursion. Observations are sorted by firm j's month $t-1$ credit rating and the regression is estimated separately for each group. The regressions are estimated jointly with QML, where the correlation between different firms' regression errors is ρ, assumed to be constant across firms and time. Noncontemporaneous innovations are assumed independent. T-statistics are in parentheses.

tion (12) is estimated with QML separately for each group. Again, the variance-covariance matrix of the residuals is the same as that in Equation (9). The results for rating categories Aaa through Baa are reported in Table 6. Those for Aaa-rated firms are based on very few observations (the sample is almost entirely composed of data for General Electric) and thus they will not be discussed further here.

As expected, for all credit rating categories, stock returns and default probabilities move in opposite directions. More importantly, the results provide strong evidence of stale bond prices. Of the total response, in a regression sense, of $h^*_{j,t}$ to lags 0–3 of firm j's stock returns, between 27% (Aa-rated firms) and 39% (Baa-rated firms) of the response occurs in the contemporaneous month. The remainder is spread out fairly evenly over the lagged months. For all rating categories (excluding Aaa), the response to each lagged stock return is statistically significant from zero at the 5% level.

The strength of the inverse relation between stock prices and instantaneous default probabilities is larger for Baa-rated firms than for either Aa-rated or A-rated firms. This result is generally consistent with models of default risk in the style of Merton (1974), which imply that the change in default risk for a given stock-price change is larger for firms that are close to the default boundary than for firms that are far away from the default boundary. The results in Table 6 indicate that a one standard deviation positive stock return in month t for a typical Aa-rated firm (which is 6.6% for this dataset) corresponds to a decline in the firm's month $t+3$ instantaneous default probability of roughly 0.00027, which corresponds to a decline in near-zero-maturity yield spreads of 1.5 b.p. [13] A one standard deviation stock return in month t for a typical Baa-rated firm (8.5% for this dataset) corresponds to a decline in the firm's month $t-3$ default probability of roughly 0.00062, or a decline in near-zero-maturity yield spreads of 3.5 b.p. This magnitude is 2.3 times the corresponding magnitude for Aa-rated firms.

7. Concluding Remarks

This article uses the extended Kalman filter to fit yields on bonds issued by individual investment-grade firms to a model of instantaneous default risk. The results are encouraging in a number of respects. First, the average error in fitting corporate bond yields is less than 10 b.p. Second, the model's parameter estimates imply that regardless of how much the financial health of a typical firm improves, yield spreads for that firm will not be driven to zero. This suggests that the model is successfully capturing the presence of a liquidity component (or more generally, a nondefault component) in yield spreads, even though the estimation procedure incorporated no information about the relative liquidities of corporate and Treasury bonds. Third, the model naturally produces an important feature of observed yield spreads; term structures of spreads for lower quality firms are more steeply sloped than are term structures of spreads for higher quality firms.

The model is by no means a complete success. There appears to be some parameter instability as firms' credit qualities change, although it is impossible to tell whether this indicates misspecification or simply randomness in parameter estimates. In addition, the model implies that the volatility of instantaneous default risk follows a square-root process, but the data indicate the presence of an additional form of persistent variation in volatility. Finally, the model is fit to observations on firms' bond yields that are themselves flawed. The yields, which are traders' indicative bid prices, appear to react very slowly to information in the firms' stock prices.

The results here can be used both as benchmarks for models of corporate bond pricing and as directions for future research. More generally they provide support for the idea that the class of instantaneous default-risk models can price default-risky claims successfully.

Appendix: Corporate Bond Pricing

The manner in which Equation (5) incorporates correlations between the default-free term structure and the default process allows default-risky zero-coupon, no recovery

[13] This computation converts $\zeta_{j,t}$ into predictions of changes in $h^*_{j,t}$ using the median standard deviation of $h^*_{j,t}$ for Aa-rated firms in Table 4 and then discounting changes in $h^*_{j,t}$ at the median κ_j also reported in Table 4.

bond prices $V_j(t, T, 0, 0)$ to be decomposed into three independent components. For each firm j, define "adjusted" riskfree rate processes $s^*_{i,t} = s_{i,t}(1 + \beta_{i,j})$, $i = 1, 2$, where under the equivalent martingale measure

$$ds^*_{i,t} = [\kappa_i \theta^*_i - (\kappa_i + \lambda_i) s^*_{i,t}]dt + \sigma^*_i \sqrt{s^*_{i,t}} d\hat{Z}_{i,t}, \tag{A1}$$

$$\theta^*_i = \theta_i(1 + \beta_{i,j}), \tag{A2}$$

$$\sigma^*_i = \sigma_i \sqrt{1 + \beta_{i,j}}. \tag{A3}$$

The $d\hat{Z}_{i,t}$ in Equation (A1) is identical to $d\hat{Z}_{i,t}$ in Equation (3b), and therefore the instantaneous innovations in the adjusted riskfree factors $s^*_{1,t}$ and $s^*_{2,t}$ are independent of each other and of the innovation in $h^*_{j,t}$. The prices of corporate zero-coupon bonds with no recovery in the event of default can be written as the following expectation under the equivalent martingale measure:

$$\begin{aligned} V_j(t, T, 0, 0) \; = \; & \exp\left[-(T - t)(\alpha_r + \alpha_j - \beta_{1,j}\overline{s_{1,t}} - \beta_{2,j}\overline{s_{2,t}}) \right] \\ & \cdot E^Q_t \exp\left[-\int_t^T s^*_{1,u} du \right] E^Q_t \exp\left[-\int_t^T s^*_{2,u} du \right] \\ & \cdot E^Q_t \exp\left[-\int_t^T h^*_{j,u} du \right] \end{aligned} \tag{A4}$$

Standard multifactor bond pricing techniques can be applied to produce closed-form solutions for Equation (A4) given the adjusted processes of Equations (A1), (A2), and (A3). See, for example, Duffie (1996a).

I use $P(t, T, c)$ to denote the price of a default-free bond that pays a coupon c at T (in addition to the principal payment) and every 6 months prior to T. Similarly, denote the price of a default-risky bond issued by firm j that promises a semiannual coupon c, but pays nothing in default, as $V_j(t, T, c, 0)$. Simple no-arbitrage arguments reveal that the price of a coupon bond is the sum of the prices of the individual coupon payments and principal payment. If coupon payments are promised at times τ_i, $i = 1, \ldots, N$, the bond prices are

$$P(t, T, c) = c \sum_{i=1}^{N} P(t, \tau_i, 0) + P(t, T, 0), \tag{A5}$$

$$V_j(t, T, c, 0) = c \sum_{i=1}^{N} V_j(t, \tau_i, 0, 0) + V_j(t, T, 0, 0). \tag{A6}$$

Finally, consider a coupon bond issued by firm j, maturing at time T, with recovery rate δ. Denote its price by $V_j(t, T, c, \delta)$. I assume that in the event of default at time τ, the bond pays off $\delta P(\tau, T, c)$. (This is equivalent to assuming the bond pays δc at each coupon payment date after default, and an additional δ at maturity.) Given this assumption, a modification of Equation (7) will hold:

$$V_j(t, T, c, \delta) = \delta P(t, T, c) + (1 - \delta) V_j(t, T, c, 0). \tag{A7}$$

Thus standard corporate bonds, which pay coupons and have recovery value in the event of default, can be priced with Equation (A7) using default-free coupon bond prices [Equation (A5)] and the prices of defaultable bonds with no recovery value [Equations (A4) and (A6)].

References

Chen, R.-R., and L. Scott, 1993, "Maximum Likelihood Estimation for a Multifactor Equilibrium Model of the Term Structure of Interest Rates," *Journal of Fixed Income*, 3, 14–31.

Chen, R.-R., and L. Scott, 1995, "Multi-factor Cox-Ingersoll-Ross Models of the Term Structure: Estimates and Tests From a Kalman Filter Model," working paper, University of Georgia.

Davidian, M., and R. J. Carroll, 1987, "Variance Function Estimation," *Journal of the American Statistical Association*, 65, 1509–1526.

Duan, J.-C., and J.-G. Simonato, 1997, "Estimating and Testing Exponential-affine Term Structure Models by Kalman Filter," working paper, McGill University.

Duffee, G. R., "The Relation Between Treasury Yields and Corporate Bond Yield Spreads," forthcoming in *Journal of Finance*.

Duffie, D., 1996a, *Dynamic Asset Pricing Theory* (2d ed.), Princeton University Press, Princeton, N.J.

Duffie, D., 1996b, "Special Repo Rates," *Journal of Finance*, 51, 493–526.

Duffie, D., and K. J. Singleton, 1995, "Modeling Term Structures of Defaultable Bonds," working paper, Stanford Graduate School of Business.

Duffie, D., and K. J. Singleton, 1997, "An Econometric Model of the Term Structure of Interest-Rate Swap Yields," *Journal of Finance*, 52, 1287–1321.

Fons, J. S., 1994, "Using Default Rates to Model the Term Structure of Credit Risk," *Financial Analysts Journal*, September/October, 25–32.

Fons, J. S., L. Carty, and J. Kaufman, 1994, "Corporate Bond Defaults and Default Rates 1970–1993," *Moody's Special Report*, January.

Geyer, A. L. J., and S. Pichler, 1996, "A State-space Approach to Estimate and Test Multi-factor Cox-Ingersoll-Ross Models of the Term Structure," working paper, Vienna University of Technology.

Grinblatt, M., 1995, "An Analytical Solution for Interest Rate Swap Spreads," Working Paper 9-94, Anderson Graduate School of Management, UCLA.

Helwege, J., and C. M. Turner, 1997, "The Slope of the Credit Yield Curve for Speculative-grade Issuers," working paper, Federal Reserve Bank of New York; forthcoming in *Journal of Finance*.

Jarrow, R. A., D. Lando, and S. M. Turnbull, 1997, "A Markov Model for the Term Structure of Credit Risk Spreads," *Review of Financial Studies*, 10, 481–523.

Jarrow, R. A., and S. M. Turnbull, 1995, "Pricing Derivatives on Financial Securities Subject to Credit Risk," *Journal of Finance*, 50, 53–86.

Jones, E. P., S. P. Mason, and E. Rosenfeld, 1984, "Contingent Claims Analysis of Corporate Capital Structures: An Empirical Investigation," *Journal of Finance*, 39, 611–625.

Kwan, S. H., 1996, "Firm-specific Information and the Correlation Between Individual Stocks and Bonds," *Journal of Financial Economics*, 40, 63–80.

Lando, D., 1994, "On Cox Processes and Credit Risky Bonds," Working Paper 1994-9, Institute of Mathematical Statistics, University of Copenhagen.

Litterman, R., and T. Iben, 1991, "Corporate Bond Valuation and the Term Structure of Credit Spreads," *Journal of Portfolio Management*, Spring, 52–64.

Longstaff, F. A., and E. S. Schwartz, 1995, "A Simple Approach to Valuing Risky Fixed and Floating Rate Debt," *Journal of Finance*, 50, 789–820.

Madan, D. B., and H. Unal, 1994, "Pricing the Risks of Default," Working Paper 94-16, Wharton School, University of Pennsylvania.

Merton, R. C., 1974, "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance*, 29, 449–470.

Nielsen, S. S., and E. I. Ronn, 1996, "The Valuation of Default Risk in Corporate Bonds and Interest Rate Swaps," working paper, University of Texas at Austin.

Pearson, N. D., and T.-S. Sun, 1994, "Exploiting the Conditional Density in Estimating the Term Structure: An Application to the Cox, Ingersoll, and Ross Model," *Journal of Finance*, 49, 1279–1304.

Sarig, O., and A. Warga, 1989, "Some Empirical Estimates of the Risk Structure of Interest Rates," *Journal of Finance*, 44, 1351–1360.

Warga, A., 1997, "A Fixed Income Data Base," working paper, University of Houston.

Zhou, C., 1997, "A Jump-diffusion Approach to Modeling Credit Risk and Valuing Defaultable Securities," Working Paper 1997-15, Federal Reserve Board.