

Is Default Event Risk Priced in Corporate Bonds?

Joost Driessen

University of Amsterdam

This article provides an empirical decomposition of the default, liquidity, and tax factors that determine expected corporate bond returns. In particular, the risk premium associated with a default event is estimated. The intensity-based model is estimated using bond price data for 104 US firms and historical default rates. Significant risk premia on common intensity factors and important tax and liquidity effects are found. These components go a long way towards explaining the level of expected corporate bond returns. Adding a positive default event risk premium helps to explain the remaining error, although this premium cannot be estimated with high statistical precision.

Compared to the extensive literature on risk premia in equity markets, relatively little is known about risk premia in corporate bond markets. Recent empirical evidence suggests that corporate bonds earn an expected excess return over default-free government bonds, even after correcting for the likelihood of default (see Section 1). In order to explain this excess return (or, equivalently, excess credit spread), existing research has looked at tax and liquidity effects, and risk premia on systematic changes in credit spreads (if no default occurs). However, a complete empirical analysis that incorporates all proposed components is lacking.

The main contribution of this article is twofold. First, an empirical decomposition of expected corporate bond returns into the several proposed determinants—interest rate and default risk premia, and tax and liquidity factors—is provided. Second, this article is, to my best knowledge, the first to estimate the risk premium associated with a default event, and to assess its importance for expected corporate bond returns.

This approach thus explicitly distinguishes the risk of credit spread changes, if no default occurs, from the risk of the default event itself. Estimation of the default event risk premium allows one to test the assumptions underlying the conditional diversification hypothesis of

I thank Frank De Jong, Siem-Jan Koopman, David Lando, Hayne Leland, Pascal Maenhout, Bertrand Melenberg, Theo Nijman, and Ilya Strebulaev for comments and discussions. I also thank Kenneth Singleton (the editor) and two anonymous referees for many helpful comments and suggestions. Finally, I thank participants at the 2001 ESSFM Gerzensee meeting, the EFA 2002 Berlin meeting, the 2002 CREDIT Conference on Credit Risk in Venice, and the Thiele Symposium at the University of Copenhagen, and seminar participants at INSEAD, Tilburg University, the Tinbergen Institute, NIB Capital Management, and ABN-AMRO Bank for their comments. Address correspondence to: Joost Driessen, Finance Group, Faculty of Economics and Econometrics, University of Amsterdam, Roetersstraat 11, 1018 WB, Amsterdam, The Netherlands, or e-mail: j.j.a.g.driessen@uva.nl.

Jarrow et al. (2001), who proved that default jump risk cannot be priced if default jumps are conditionally independent across an infinite number of firms. A positive jump risk premium would thus indicate that default jumps are not conditionally independent across firms or that not enough corporate bonds are traded to fully diversify the default jump risk.

The Duffie and Singleton (1999) framework is used to specify the model. In case of default a downward jump in the bond price occurs, equal to a (constant) loss rate times the bond price just before default. The product of the risk-neutral jump intensity and the loss rate equals the instantaneous credit spread. Each firm's default intensity is modeled as a function of latent common factors and a latent firm-specific factor. This extends the results of Duffee (1999), who estimated a separate model for each firm. All factors follow square-root diffusion processes. In line with the results of Longstaff and Schwartz (1995) and Duffee (1999), the model also allows for correlation between credit spreads and default-free interest rates, which are modeled using a two-factor "essentially affine" model [Duffee (2002)]. The model also corrects for tax differences between corporate and government bonds, and a liquidity factor that is based on the corporate bond age is included. Finally, the ratio of the risk-neutral and actual jump intensities defines the jump risk premium.

In sum, the model allows for six different components of expected excess corporate bond returns. Two components are risk premia on, respectively, market-wide and firm-specific changes in credit spreads. A third component is due to the dependence of credit spreads on default-free interest rates, and a fourth component is due to a risk premium on the default jump. The remaining two components are tax and liquidity effects.

The estimation methodology consists of two steps. First, the Kalman filter maximum likelihood is used to estimate the affine factor model for risk-neutral intensities. This step uses weekly US corporate bond price data for 592 bonds of 104 firms, from 1991 to 2000. The results show that the market-wide spread risk, represented by movements in two common factors, is priced, whereas the firm-specific factor risk is not. A negative relationship is found between credit spreads and the default-free term structure. The liquidity factor, which represents the liquidity difference between recently issued and old bonds, is another important determinant of credit spreads. In particular, a downward sloping term structure of liquidity spreads is found. Finally, tax effects also explain part of the level of credit spreads.

Next, the implications of this estimated model for actual default probabilities, excluding a default jump risk premium, are studied. In this case the model overestimates observed default rates, and underestimates expected excess corporate bond returns and credit spreads. This "credit spread puzzle" has also been reported by others (see Section 1).

Next, a second estimation step is performed, in which the default jump risk premium is estimated using data on historical default rates. The estimated premium turns out to be positive, and it leads to a better fit of observed default rates. The standard error of the estimated premium is obtained by incorporating the sampling error in historical default rates and in the estimates of the factor model parameters. With a t -ratio of 1.2, the statistical evidence for the existence of a default jump risk premium is inconclusive. In other words, once common factor risk premia and tax and liquidity effects are included, the statistical evidence for the existence of a credit spread puzzle is not very strong.

The economic importance of all factors is assessed by decomposing the expected corporate bond return into the different components. This analysis shows that, economically, the jump risk premium mainly matters for BBB-rated bonds, explaining about 31 basis points of the expected annual return on a 10-year bond. For this rating category, the tax effect and common factor risk are the two other important determinants. For AA-rated and A-rated bonds the most important determinants are tax and liquidity effects, and common factor risk. The default jump premium contributes, respectively, 6 and 12 basis points to the expected return of AA-rated and A-rated bonds.

The rest of the article is organized as follows. Section 1 discusses the related literature, Section 2 introduces the model, and Section 3 describes the corporate bond data. Section 4 discusses the estimation methodology and the results for the risk-neutral intensity model. Section 5 describes the estimation of the default jump risk premium, followed by a presentation of the results and robustness checks. Section 6 concludes.

1. Contribution to Existing Literature

This article is most closely related to a set of articles that also evaluate models using both corporate bond price data and historical default rates [Elton et al. (2001), Jarrow et al. (2001), and Huang and Huang (2002)]. Leland (2002) pointed out the importance of evaluating corporate bond pricing models using observed default rates, besides using corporate bond price data.

Jarrow et al. (2001) performed a first analysis of conditional diversification, but they did not estimate the default event risk premium. They use the estimates of Duffee's (1999) model to compare the model-implied default probabilities (in the absence of a default jump risk premium) with the default probabilities that are implied by an annual Markov migration model. The results on the event risk premium obtained in this article differ from the results of Jarrow et al. (2001) (see Section 5.2). In particular, it is shown that the use of the Markov migration model leads to downward biased estimates of the default jump risk premium.

Elton et al. (2001) show that expected default losses and tax effects cannot explain the observed level of credit spreads. They explain part of this fitting error by relating corporate bond returns to equity market factors. Although this article also finds that common movements in corporate bond prices carry a risk premium, there are several differences between the study by Elton et al. and this article. First, this article uses a latent factor model instead of the Fama-French factor model, because Collin-Dufresne et al. (2001) show that observable financial and economic variables (including equity returns) cannot explain the correlation of credit spread changes across firms. Moreover, Elton et al. do not incorporate a liquidity factor and a default event risk premium.

Huang and Huang (2002) fit structural firm value models to, amongst others, historical default rates and the equity premium, and compare model-implied and observed credit spreads. They find that all models imply too low investment grade credit spreads. Most of these models are based on a diffusion process for changes in the firm value. The risk premia generated by these models can be interpreted as risk premia on common changes in firm values, or, equivalently, credit spreads. It is then interesting to see that Huang and Huang explain about 30% of the BBB-credit spread, which is similar to the explanatory power of the risk premia of the common spread factors in this article. Huang and Huang also consider a model that incorporates systematic (and priced) jumps in the firm value. Although these priced jumps do not necessarily cause a default event, they may capture a similar effect as the priced default jumps in this model. Huang and Huang find that the impact of jumps depends strongly on the parameter values for jump intensity and size, but that it is unlikely that the incorporation of jumps fully explains the credit spread puzzle. I find that tax and liquidity effects (not incorporated by Huang and Huang) are important determinants of credit spreads, so that it is indeed unlikely that default risk alone can explain credit spread levels. An important difference between Huang and Huang and this article is that here the model is estimated using firm-level data, while Huang and Huang calibrate structural models to aggregate data without analyzing the statistical accuracy of their results.

Delianedis and Geske (2001) study only the level of credit spreads, and did not use historical default rates. They use a firm value framework to assess the influence of several factors on the level of credit spreads, including tax, jump, and liquidity factors. Although they do not empirically decompose the credit spread into all these factors, they argue that priced jumps in firm value cannot solely explain the high level of credit spreads, in line with the results of Huang and Huang and this article.

Collin-Dufresne et al. (2003) propose a theoretical model where a large upward jump in a firm's credit spread causes a moderate market-wide

jump in the credit spreads of other firms (a contagion effect). They distinguish the direct jump risk premium, which is similar to the default event risk premium in this article, from the contagion risk premium. The market-wide jumps of Collin-Dufresne et al. are different from the common diffusion factors for credit spreads in this article, but empirically they may capture similar effects.¹ Collin-Dufresne et al. perform a calibration to obtain an indication of the size of the jump and contagion risk premia, which is compared to the results in Section 5.2.

Janosi et al. (2001) estimate an intensity-based model in which common factors influence the intensities. They also incorporate a reduced-form liquidity correction, assuming that the relationship between liquidity and variables such as the spot rate, equity market return, and volatility, is different from the relationship between the intensities and these variables. However, they do not incorporate tax effects or a default event risk premium. Their focus is on the bond pricing performance of the model, and their common factor model gives a good fit of corporate bond prices, in line with the results of this article.

Yu (2002) also provides a decomposition of corporate bond returns using the intensity-based framework of Duffie and Singleton (1999), but he did not estimate the size of the components.

2. Model for Defaultable Bond Prices

2.1 Model setup

For US Treasury bonds (assumed to be default-free) I use a two-factor “essentially affine” model [Duffee (2002)], in which the instantaneous default-free short rate r_t is a function of two factors, $r_t = \delta_0 + \delta' X_t$, where $X_t = (X_{1,t}, X_{2,t})'$ and $\delta = (\delta_1, \delta_2)'$. Following the notation of Dai and Singleton (2000), the $A_1(2)$ model is used. The model generates time-varying risk premia and correlation between the factors, which is important to fit interest rate data [Dai and Singleton (2000)]. The factors follow an affine process under the actual probability measure P

$$\begin{bmatrix} dX_{1,t} \\ dX_{2,t} \end{bmatrix} = \begin{bmatrix} \kappa_{11} & 0 \\ \kappa_{21} & k_{22} \end{bmatrix} \begin{bmatrix} \theta_1 - X_{1,t} \\ -X_{2,t} \end{bmatrix} dt + \begin{bmatrix} \sqrt{X_{1,t}} & 0 \\ 0 & \sqrt{1 + \beta_{21} X_{1,t}} \end{bmatrix} dW_t. \quad (1)$$

W_t is a two-dimensional vector of independent Brownian motions. Bond risk premia are modeled following the specification proposed by Duffee (2002). The Brownian motion vector $\hat{W}_t$ under a risk-neutral probability

¹ It may be econometrically hard to distinguish between the effects of a market-wide jump process and common diffusion factors using discretely observed credit spread data. A recent example of work in this direction is Bierens et al. (2003).

measure Q satisfies $d\hat{W}_t = \Lambda_t dt + dW_t$, where

$$\Lambda_t = \begin{bmatrix} \sqrt{X_{1,t}} & 0 \\ 0 & \sqrt{1 + \beta_{21} X_{1,t}} \end{bmatrix} \begin{bmatrix} \lambda_{11} \\ \lambda_{12} \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & (1 + \beta_{21} X_{1,t})^{-1/2} \end{bmatrix} \begin{bmatrix} 0 & 0 \\ \lambda_{2(21)} & \lambda_{2(22)} \end{bmatrix} \begin{bmatrix} X_{1,t} \\ X_{2,t} \end{bmatrix}. \quad (2)$$

Duffee (2002) shows that this model leads to the well-known exponential-affine pricing formula for default-free bonds.

As in Duffee and Singleton (1999), Madan and Unal (1998), and Jarrow and Turnbull (1995), a default event is modeled as the first jump of a conditional Poisson process. The stochastic intensity of this jump process at time t under P is denoted by $h_{j,t}^P$, for firm j , $j = 1, \dots, N$. In case of a default event at time t , there is a downward jump in the bond price equal to $L_{j,t}$ times the price of the bond just before the default event. Duffee and Singleton (1999) call this the recovery of market value assumption. As in Duffee (1999) and Elton et al. (2001), this loss rate is assumed to be constant at 56%. Below, it is shown that this assumption is only used for the estimation of the default event risk premium, and not for the estimation of the risk premia of credit spread changes.

Assuming the absence of arbitrage opportunities guarantees the existence of an equivalent martingale measure Q . As noted by Jarrow et al. (2001), the intensity $h_{j,t}^Q$ under this measure is related to the P -intensity through the risk premium parameter μ

$$h_{j,t}^Q = \mu h_{j,t}^P. \quad (3)$$

If the risk associated with default events is priced, the parameter μ will exceed 1. Although this risk premium parameter may vary over time, it is assumed to be constant for simplicity.

Duffee and Singleton (1999) show that the time t price $V_j(t, T)$ of a defaultable zero-coupon bond (that has not yet defaulted), issued by firm j and maturing at time T , is given by

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

with instantaneous spread $s_{j,t} \equiv h_{j,t}^Q L$ and where E_t^Q denotes the Q -expectation conditional upon the information set at time t . This shows that it suffices to model the instantaneous spread $s_{j,t}$ to price defaultable bonds. As mentioned above, this also implies that for the modeling of credit spreads and credit spread risk premia, the assumption of a constant loss rate is innocuous.

In order to apply formula (4) to the data on coupon-paying corporate bond prices, this pricing formula is modified in two ways. First, coupons on corporate bonds are subject to state taxes, while government bond coupons are not [Elton et al. (2001)]. In Appendix A it is shown how this tax effect can be incorporated in the intensity-based pricing model. Elton et al. (2001) estimated the effective tax rate to be 4.875%. For most results, this tax rate is used. Second, corporate bonds are less liquid than government bonds (as discussed in Section 1), which motivates the inclusion of a common liquidity component $(\beta_{0j}^l + \beta_{1j}^l l_t)$ in the instantaneous credit spread, so that $s_{j,t}^Q \equiv h_{j,t}^Q L + (\beta_{0j}^l + \beta_{1j}^l l_t)$. The liquidity factor l_t is assumed to follow a standard square-root diffusion process under P and Q . In Section 4.2 we discuss the interpretation and estimation of this liquidity component.

Given the evidence for systematic changes in credit spreads across firms [see Collin-Dufresne et al. (2001) and Elton et al. (2001)], the risk-neutral intensity of firm j is modeled as a function of K common factors $F_{i,t}$, $i = 1, \dots, K$, and a firm-specific factor $G_{j,t}$, plus two terms that allow for correlation between spreads and default-free rates

$$h_{j,t}^Q L = \gamma_{0j} + \sum_{i=1}^K \gamma_{ij} (F_{i,t} - \theta_i^F) + (G_{j,t} - \theta_j^G) + \alpha_{1j} (X_{1,t} - \theta_1) + \alpha_{2j} X_{2,t}. \quad (5)$$

A latent factor model is used since Collin-Dufresne et al. (2001) show that financial and economic variables cannot explain the correlation structure of credit spreads across firms. The K common factors and N firm-specific factors follow independent square-root processes under P

$$\begin{aligned} dF_{i,t} &= \kappa_i^F (\theta_i^F - F_{i,t}) dt + \sigma_i^F \sqrt{F_{i,t}} dW_{i,t}^F, \quad i = 1, \dots, K \\ dG_{j,t} &= \kappa_j^G (\theta_j^G - G_{j,t}) dt + \sigma_j^G \sqrt{G_{j,t}} dW_{j,t}^G, \quad j = 1, \dots, N. \end{aligned} \quad (6)$$

Here, the κ -parameters are mean-reversion parameters, the θ -parameters represent the unconditional factor means, and the σ -parameters can be interpreted as volatility parameters.² Both for the common and firm-specific factors, it is assumed that the market price of factor risk depends on the factor level. For example, for each common factor it is assumed that $d\hat{W}_{i,t}^F = dW_{i,t}^F + (\lambda_i^F / \sigma_i^F) \sqrt{F_{i,t}} dt$, where $\hat{W}_{i,t}^F$ is a Brownian motion under Q . For the liquidity factor l_t a similar assumption for the liquidity risk premium λ^l is made.

Given the affine processes for the factors under Q , Equation (4) implies that corporate bond prices are exponential-affine functions of the factors [see Duffie and Kan (1996)].

² Not all parameters in the processes in Equation (6) can be identified. In Appendix B it is shown that the identification problem can be solved by normalizing the means of the factors θ_i^F , $i = 1, \dots, K$.

2.2 Expected bond returns and conditional diversification

Let us start with the default-free expected bond returns. Let $P(t, T)$ denote the time t default-free bond price for maturity date T . Applying Ito's lemma to the well-known exponential-affine bond price expression of Duffie and Kan (1996) gives

$$E_t^P \left[\frac{dP(t, T)}{P(t, T)} \right] = r_t dt + \tilde{A}(\delta, X_t, T - t) dt \quad (7)$$

with

$$\begin{aligned} \tilde{A}(\delta, X_t, T - t) \equiv & -(T - t)[A_1(\delta, T - t)A_2(\delta, T - t)] \\ & \times \left(\begin{bmatrix} 0 \\ \lambda_{12} \end{bmatrix} + \begin{bmatrix} \lambda_{11} & 0 \\ \lambda_{12}\beta_{12} & 0 \end{bmatrix} X_t + \begin{bmatrix} 0 & 0 \\ \lambda_{2(21)} & \lambda_{2(22)} \end{bmatrix} X_t \right). \end{aligned} \quad (8)$$

The functions $A_1(.,.)$ and $A_2(.,.)$ can easily be solved for numerically. The notation used in this article, incorporates explicitly the dependence of the functions $A_1(.,.)$ and $A_2(.,.)$ on the parameter vector $\delta = (\delta_1, \delta_2)'$. This notation will be useful below.

For corporate bond returns, the expression is slightly more complicated, because one has to incorporate the influence of a default event on the expected return. Using the results of Yu (2002), Appendix C derives the following expression for the instantaneous expected return on a corporate discount bond, in excess of a government bond with the same maturity

$$\begin{aligned} & \left[-\sum_{i=1}^K (T - t) B_{ij}(T - t) \lambda_i^F F_{i,t} - (T - t) C_j(T - t) \lambda_j^G G_{j,t} \right. \\ & \quad + \tilde{A}(\alpha_j + \delta, X_t, T - t) - \tilde{A}(\delta, X_t, T - t) + (\mu - 1) h_{j,t}^P L \\ & \quad \left. + (\beta_{0j}^l + \beta_{1j}^l l_t - (T - t) D_j(T - t) \lambda^l l_t) \right] dt. \end{aligned} \quad (9)$$

where $\alpha_j = (\alpha_{1j}, \alpha_{2j})'$ and where the functions $B_{ij}(\cdot)$, $C_j(\cdot)$, and $D_j(\cdot)$ depend on the model parameters [see, e.g., Pearson and Sun (1994) for explicit expressions for these loading functions in square-root models]. In total, the model generates expected excess corporate bond returns in six ways. First, because the risk of common changes in credit spreads is priced. Second, via a risk premium on firm-specific credit spread changes. Third, through the dependence of credit spreads on default-free factors. Fourth, expected bond returns are positively related to the jump risk premium $(\mu - 1)$. Fifth, due to the liquidity difference between corporate and government bonds. Finally, for coupon-paying bonds, the tax difference between corporate and government bonds also explains part of the pre-tax excess corporate bond return.

At this stage it is appropriate to discuss the “asymptotic equivalence” results of Jarrow et al. (2001). In an intensity-based framework, they show that the P -intensity is approximately equal to the Q -intensity ($\mu \simeq 1$), if (1) default processes are independent, conditional on the path of default intensities, and (2) there is a countable infinite number of firms in the economy. Jarrow et al. refer to this as a case of “diversifiable default risk”. Intuitively, given these two assumptions, default jumps can be (approximately) diversified away and are therefore not priced.

By estimating the default jump risk premium, it is possible to test the assumptions underlying this “conditional diversification” hypothesis. Jarrow et al. mention two reasons why default jump risk could be priced. First, there may be a (small) possibility that some firms default simultaneously. Second, in an economy with a finite number of bonds a positive default jump risk premium may exist even if default jumps are conditionally independent.

Because it is not necessary to specify whether default events are conditionally independent or not in order to price corporate bonds that have not yet defaulted [Equation (4)], both explanations for a positive default jump risk premium are in line with the model for default events used in this study. One interpretation of the model is that the common factors drive intensities, with independent defaults conditional upon this intensity process. The model can however also allow for simultaneous default, if a common intensity factor $F_{i,t}$ describes the arrival rate of joint default events. This also implies that it is not possible to distinguish empirically between these two interpretations of the model, because of the use, below, of data on bond prices of firms that did not default during the sample period.

3. Description of the Dataset

Data on mid-quotes for US-dollar corporate bond prices and bond characteristics are taken from the Bloomberg Corporate Bonds Database. Weekly data from February 22, 1991, until February 18, 2000, were collected. This study is restricted to the set of 161 firms analyzed by Duffee (1999). A firm is included in this analysis only if there are data on at least two corporate bonds for at least 100 weeks, which leaves 104 of the 161 firms. This analysis only uses bonds with constant, semi-annual coupon payments that do not contain any put or call options, or sinking fund provisions. As in Duffee (1999), observations on bond prices with less than one year to maturity were dropped. Also, if the maturity difference of two bonds is smaller than six months, only the most recently issued bond is retained. More than 80% of the remaining bonds is senior unsecured. Other bonds were included only if a bond has the same rating as the senior unsecured bonds. At the end of the sample period, all 104 firms are rated

Table 1
Summary statistics for corporate bond data

Firm-level statistic	Across 104 firms		
	Minimum	Median	Maximum
Weeks of data	100	378	470
Mean number of fitted bonds per week	2.0	3.0	8.0
Mean years to maturity of fitted bonds	2.6	7.7	22.9
Minimum years to maturity of fitted bonds	1.0	1.0	11.8
Maximum years to maturity of fitted bonds	4.7	15.0	50.0
Mean coupon of fitted bonds	5.8	7.6	10.1

The table reports summary statistics on weekly observations for corporate bond prices from February 22, 1991, until February 18, 2000, for 592 bonds of 104 firms. The row “Weeks of data” contains the number of weeks for which at least two bond prices are observed for a given firm. “Mean number of fitted bonds” contains the mean number of bonds fitted per week, conditional upon two bond prices observed during this week.

investment-grade: 2 AAA-rated firms, 13 AA-rated firms, 58 A-rated firms, and 31 BBB-rated firms.

Table 1 contains information on the bond data. For the median firm, on average three different bonds are used to estimate the model, and at 378 of the 470 weeks in the data at least two bond prices are observed for this median firm. Besides corporate bond price data, this study also uses Bloomberg data on the six-month US Treasury bill, and the most recently issued US Treasury bonds, for the maturities at 2, 3, 5, 7, 10, and 30 years.

To further analyze the corporate bond price data, coupon spreads are calculated. The coupon spread is defined as the difference between a corporate bond yield and the yield of a default-free bond with the same coupon and maturity. This study uses the estimation results for the two-factor affine model in Equation (1), which will be discussed in Section 4, to calculate the associated yield-to-maturities of default-free coupon bonds.³ Figure 1 plots the time series of the coupon spreads, averaged within each rating category. The graph shows that from 1991 to 1998 spreads have declined. Due to the Russia/LTCM crisis in the fall of 1998, spreads increased dramatically, and have remained high since. Figure 1 also shows that there is considerable correlation between the spreads of the different rating categories, which motivates the common factor model. Unreported results show that, on average, high-rated firms have a low and slowly increasing spread term structure, whereas low-rated firms have higher and more steeply increasing spread term structures, in line with the summary statistics in Duffee (1999).

³ Some bond price observations are very likely incorrect. Observations for which the coupon spread is above 400 basis points or below -50 basis points, and observations that relate to a coupon spread movement of more than 100 basis points in one week are eliminated. The “middle” observation of observations for which the coupon spread moves more than 50 basis points in one week, and again more than 50 basis points in the opposite direction in the next week are also deleted. Thus, 616 of the 140,389 bond price observations are eliminated.




Figure 1

Time series of the coupon spreads

For each week, coupon spreads of 592 bonds of 104 firms are averaged within each rating category. The graph depicts the resulting time series.

4. Estimation of the Factor Model for Risk-Neutral Intensities

The setup of the model is such that it can be estimated in two steps. First, this section describes the estimation methodology for the model for the risk-neutral intensity $h_{j,t}^Q$. The estimation of the default jump risk premium μ is described in Section 5.

4.1 Estimation methodology

Similar to Duffee (1999), this study uses quasi maximum likelihood (QML) based on the Kalman filter to estimate the model for risk-neutral intensities in Equations (5) and (6). Appendix D provides details on the Kalman filter QML. As in Duffee (1999) this estimation directly uses the yield-to-maturities of coupon-paying Treasury and corporate bonds to estimate the parameters. This avoids using an ad-hoc smoothing method to calculate zero-coupon interest rates.

In principle, a joint estimation of all parameters is most efficient. However, since the number of parameters is large, the estimation is performed in four steps. First, Kalman filter QML is used to estimate the two-factor affine model in Equation (1) for the default-free term

structure, using data on Treasury-bill rates and Treasury bond yields. The second step of the estimation procedure involves estimation of the parameters in the square-root process for the liquidity factor l_t . This step is discussed in Section 4.2 and also provides a time series of liquidity factor values $l_1, \dots, l_T$.

Using the model for default-free Treasury rates, the coupon spreads are calculated. These spreads are used in step three, which involves the estimation of the common factors parameters and the liquidity parameters β_{0j} and β_{1j} , again using Kalman filter QML. The constant terms γ_{0j} in Equation (5) are also estimated in this step. To reduce the number of parameters to be estimated in the third step, the firm-specific parameters $\alpha_{1j}, \alpha_{2j}, \beta_{0j}, \beta_{1j}, \dots, \gamma_{Kj}$ are restricted to be constant across firms that have the same credit rating at the end of the sample period.⁴

In step four, the residual fitting errors for the coupon spreads from step three are used to estimate the parameters of the firm-specific factors $G_{j,t}$, $j = 1, \dots, 104$. These parameters are estimated for each firm separately, using Kalman filter QML. Finally, as described in Appendix D, the third and fourth step of this estimation procedure are repeated, where in the third step the i.i.d. measurement error assumption is replaced by the structure implied by the estimated firm-specific factor processes and the measurement error structure that was assumed in this fourth step. Thus, this method explicitly incorporates the presence of firm-specific factors when estimating the common factor processes. On analyzing whether applying another iteration leads to important changes in the parameter estimates, it was found that this is not the case.

In each estimation step standard errors and t -ratios for the parameter estimates are calculated [correcting for heteroscedasticity using the White (1982) covariance matrix], assuming that the parameters that are estimated in previous steps are estimated without error. In principle, it is possible to calculate the standard errors by taking into account the previous steps, for instance, by means of bootstrapping, but this is excessively time-consuming.

For all square-root processes in the model, the parameters are estimated given the Feller condition. For example, for the common factors, the restriction $2\kappa_i^F \theta_i^F > (\sigma_i^F)^2$, $i = 1, \dots, K$ is imposed. This restriction turns out to be not binding for the common factors. For the firm-specific factors, the restriction is binding for 7 out of the 104 firms.

4.2 Estimation of the liquidity factor

This subsection describes the construction of the liquidity factor l_t and the estimation of the β -parameters in the liquidity spread specification

⁴ Since there are only two AAA-rated firms, these two firms are treated as AA-rated firms.

$(\beta_{0j}^l + \beta_{1j}^l I_t)$. To start with, the panel dataset of coupon spreads of all firms is regressed on two liquidity proxies, the amount issued, and the age of the bond. For the rating categories, the bond maturity and dummy variables are included as control variables. To allow for maturity effects of liquidity, both the amount issued and the bond age with maturity are allowed to interact with maturity. Table 2 shows that bond age is a significant determinant of credit spreads, while the amount issued is not statistically significant. The age-effect decreases as bond maturity increases. Chakravarty and Sarkar (1999), Elton et al. (2002), and Houweling et al. (2003) also find that bond age influences credit spreads.

Based on these regression results, the bond age is used as a liquidity indicator. Thus, four portfolios of all corporate bonds are formed in this dataset. The first two (equally weighted) portfolios contain, for a given week, all bonds with bond age less than three years, with maturities, respectively, smaller and larger than five years. The other two portfolios contain bonds with bond age above three years, again with maturities, respectively, smaller and larger than five years. These latter two portfolios are not equally weighted. Instead, the weights are chosen to match the duration and the rating distribution of the first two portfolios. By calculating the difference between the average yield of the low-age portfolios

Table 2
Estimates of the liquidity premium

	Estimate	Standard error
(a) Regression of the coupon spreads on the liquidity proxies		
log(amount issued)	-0.264	0.218
log(amount issued) $\times$ maturity	-0.031	0.028
Age	1.429	0.451
age $\times$ maturity	-0.023	0.017
(b) Estimation results of the one-factor square-root model for liquidity spread		
κ^1	1.638	0.432
λ^1	1.009	0.361
σ^1	0.059	0.043
θ^1	20.58 bp	11.38 bp
(c) Factor model for credit spreads: estimates for the liquidity constant term and slope		
$\beta_{0,AA}^1$	7.48 bp	
$\beta_{0,A}^1$	6.84 bp	
$\beta_{0,BBB}^1$	7.14 bp	
$\beta_{1,AA}^1$	0.380	0.273
$\beta_{1,A}^1$	0.442	0.200
$\beta_{1,BBB}^1$	0.419	0.282

The table contains estimates for the liquidity factor in the benchmark model (Section 4.2). (a) Estimates of a regression of all coupon spreads [in basis points (bp)] on the four listed explanatory variables, bond maturity, and dummies for each rating category. The maturity and rating dummy coefficients are not shown. For each bond, standard errors are corrected for heteroscedasticity and first-order autocorrelation using the Newey and West (1987) covariance matrix. (b) Estimates of a one-factor square-root model for the spread between low-age and high-age bonds, obtained using Kalman filter QML. (c) Estimates for the constant term and slope coefficient in the liquidity model. The estimation strategy is described in Sections 4.1 and 4.2.




Figure 2

Time series of the liquidity spreads

Using Kalman filter QML, a one-factor square-root model is estimated for the spread between high-age and low-age bonds (see Section 4.2). The graph depicts the filtered factor values for this liquidity spread and the average coupon spread of all bonds for each week.

versus the high-age portfolios, it is possible to obtain, for each week, both a short-maturity and long-maturity liquidity spread.

Next, a one-factor square-root term structure model [similar to Equation (6)] is estimated for the liquidity factor l_t using the time series of these two liquidity spreads. Estimation is again performed using Kalman filter QML.⁵ Results are given in Table 2*b*. The unconditional expectation of the instantaneous liquidity spread equals 20.6 basis points. The risk premium parameter is actually positive, which implies that long-maturity bonds have a lower liquidity spread than short-maturity bonds. This is consistent with the regression results in Table 2*a*. The downward sloping liquidity term structure is also in line with the results of Janosi et al. (2002). Figure 2 plots the time series of the liquidity spread l_t (obtained using the Kalman filter), and the time series of the average corporate–government coupon spread. The graph shows a positive relationship

⁵ To simplify the estimation procedure, the liquidity spreads of the portfolios are treated as spreads on zero-coupon bonds with maturities equal to the average durations of the portfolios (3.2 and 8.3 years, respectively).

between liquidity and credit spreads. In particular, liquidity and credit spreads are both high in the beginning of the sample period and during the Russia/LTCM crisis. This is in line with the liquidity time series of Liu et al. (2001). Also, there seems to be a downward trend in the liquidity spread, which has also been documented by Chen and Wei (2002) and Janosi et al. (2002). Based on these results, it seems safe to assume that the liquidity factor is related to the time variation in market-wide liquidity, thus allowing credit spreads to load on l_t through the parameter β_{lj}^l . This slope parameter is estimated in step three of the procedure described in Section 4.1.

The liquidity model also contains a constant liquidity term β_{0j}^l , which allows one to fit the average liquidity spread. This term can be estimated separately, since the model in Equation (5) already contains a constant term γ_{0j} . Effectively, part of this constant term is interpreted as a liquidity component. The average liquidity spread is estimated using the regression results in Table 2. It is assumed that the liquidity of a bond that has just been issued with a large amount issued is similar to the liquidity of government bonds. For this bond, the age is set equal to zero and the amount issued equal to the 95th percentile of the empirical distribution of the amount issued. This is compared with the average bond, for which the amount issued is about one fifth of the 95th percentile and for which the bond age is 9.6 years. At the average maturity of 8.2 years, the regression results in Table 2a imply a yield difference of 12.7 basis points between these two bonds. For each rating category and given the estimates for all other parameters in the risk-neutral intensity model, the parameter β_{0j}^l is chosen such that the unconditional expectation of the total liquidity spread at 8.2 years maturity equals 12.7 basis points [for a bond with the coupon rate equal to the average coupon rate for the median firm (7.6%)].

4.3 Estimation results for the risk-neutral intensity model

Table 3 presents the estimation results for the default-free part of this model. The estimates for the interest rate risk premia are such that government bonds earn, on average, a positive return in excess of the risk-free short rate [Equation (8)]. The fit on the Treasury yields is reasonable. Across all maturities, the mean absolute yield error is 8.05 basis points.

Using a likelihood ratio test⁶ results in a model with two common factors for the intensity model in Equation (5). The estimates, reported in Table 4, show that the market prices of risk of both factors are negative and jointly statistically significant, indicating that corporate bond

⁶ For this likelihood ratio test, the non-normality of the factor changes is neglected and a normal approximation is used for these changes.

Table 3
Kalman filter QML estimates of the two-factor essentially affine model for default-free rates

Parameter	Estimate	Parameter	Estimate
κ_{11}	0.012 (0.007)	δ_2	0.004 (0.003)
κ_{21}	-0.182 (0.235)	β_{21}	0.871 (0.642)
κ_{22}	0.619 (0.204)	λ_{11}	-0.052 (0.041)
θ_1	2.012 (1.292)	λ_{12}	-3.481 (0.981)
δ_0	0.038 (0.002)	$\lambda_{2(21)}$	-0.326 (0.682)
δ_1	0.002 (0.001)	$\lambda_{2(22)}$	1.444 (1.092)

Using Kalman filter QML (see Appendix D), the two-factor essentially affine model in Equation (1) is estimated using weekly data on the six-month Treasury-bill rate and Treasury bond yields with maturities of 2, 3, 5, 7, 10, and 30 years. It is assumed that all interest rates are observed with i.i.d. measurement errors. Standard errors (in brackets) are corrected for heteroscedasticity using the White (1982) covariance matrix.

Table 4
Kalman factor QML estimates of the common factor model for intensities

Factor i	κ_i^F	λ_i^F	σ_i^F
Factor 1	0.049 (0.067)	-0.058 (0.024)	0.014 (0.007)
Factor 2	0.629 (0.151)	-0.162 (0.111)	0.054 (0.012)
	AA rating	A rating	BBB rating
$\gamma_{0,R}$	12.39 bp (14.62 bp)	26.28 bp (15.18 bp)	54.31 bp (21.66 bp)
$\gamma_{1,R}$	0.590 (0.152)	0.853 (0.177)	1.152 (0.200)
$\gamma_{2,R}$	0.191 (0.075)	0.359 (0.060)	0.521 (0.125)
$\delta_1\alpha_{1,R}$	-0.026 (0.006)	-0.057 (0.007)	-0.061 (0.010)
$\delta_2\alpha_{2,R}$	-0.006 (0.017)	0.001 (0.017)	-0.022 (0.018)

Using Kalman filter QML, the benchmark model with two common factors in Equations (5) and (6) is estimated (see Appendix D). Standard errors are calculated using the White (1982) heteroscedasticity-consistent covariance matrix. The parameter θ_i^F is normalized to 50 basis points for both factors.

investors demand a positive return in excess of default-free bond returns, to be compensated for the risk associated with common spread movements. This evidence for common factors in risk-neutral intensities is in line with the results of Das et al. (2002), who provide evidence for common factors in actual default probabilities across firms.

For both factors, the estimates for the factor loadings γ_{ij} are positive for all three rating categories (AA, A, and BBB). Thus, both factors influence credit spreads in the same direction for all firms. The first factor has a loading function $B_{1j}(T-t)$ in Equation (9) that is increasing with maturity ($T-t$), whereas the loading function $B_{2j}(T-t)$ is decreasing with maturity, so that the second factor mostly influences short-maturity spreads. For both factors, low-rated firms are more sensitive to the common factors than high-rated firms. The explanation for this result is twofold. First, spread term structures are more steeply increasing for lower ratings, and, second, spreads are more volatile for low-rated firms. Indeed, a higher value for γ_{ij} implies both steeper spread term structures and more volatile spreads. Finally, in line with results of Longstaff and Schwartz (1995) and Duffee (1999), there is a negative relationship between spreads

Table 5
Firm-specific factor parameter estimates

	First quartile	Median	Third quartile
Estimate κ_j^G	0.029	0.049	0.092
Estimate λ_j^G	-0.064	-0.004	0.027
Estimate σ_j^G	0.004	0.008	0.015
Estimate θ_j^G	11.59 bp	22.85 bp	39.32 bp
$K=0$: Estimate λ_j^G	-0.147	-0.108	-0.053

The table reports quartiles of parameter estimates for the firm-specific factors in the model with two common factors ($K=2$) in Equations (5) and (6). The last row reports estimates for the market price of risk in a model with $K=0$. Estimates are obtained using the QML based on Kalman filter (see Appendix D).

and default-free rates for firms in all three rating categories. The explained variation is however quite small.

Table 5 presents the estimation results for the parameters in the firm-specific factor processes. Most strikingly, the market price of the risk associated with movements in the firm-specific factors is close to zero for the median firm. For comparison, a model without common factors is also estimated [$K=0$ in Equation (5)]. This model, with firm-specific factors only, is similar to the model of Duffee (1999). In this case, the estimate for the market risk price, reported in Table 5, is negative for almost all firms, and much larger in absolute size relative to the model with $K=2$.⁷ Thus, after correcting for market-wide spread risk by including two common factors, the remaining firm-specific movements in spreads are hardly priced. To verify that the firm-specific factors are really firm specific, the cross-firm correlations of weekly changes in these firm-specific factors are calculated. The average of these cross-firm correlations is 0.076, which is much lower than 0.348, the average cross-firm correlation that is found for the model without two common factors ($K=0$).⁸

The liquidity parameters β_{0j} and β_{1j} are reported in Table 2c. The results show a positive dependence of credit spreads on the liquidity factor l_t , in line with the graph in Figure 2. In order to match the average liquidity spread of 12.7 basis points, the constant terms are also positive.

Finally, the fit of the model on corporate bond yields is assessed. The model with $K=2$ has a mean absolute yield error of 8.63 basis points for the median firm, which is in the same order of magnitude as the fit on the Treasury bonds. The firm with the worst fit has a mean absolute yield error of 17.41 basis points, which still seems reasonable. The time-series R^2 for the changes in credit spreads is also calculated. To compare with the

⁷ Estimates for the other parameters in the model with $K=0$ are qualitatively similar to the estimates reported by Duffee (1999).

⁸ A three-factor model has also been estimated. For this model the average cross-firm correlation of the firm-specific factors equals 0.049.

common factor model of Elton et al. (2001), this model excludes the firm-specific factors for this calculation and focuses on the credit spread changes of the equally weighted portfolios across ratings and one-year maturity intervals. The average R^2 across all portfolios equals 38.7%, which is much larger than the average R^2 (about 17%) of the factor model of Elton et al. (2001). For a final evaluation of the model fit, the 25 weeks that have the largest absolute movements in the coupon spreads (averaged across all firms) are selected. Then, the fit on corporate bond yields at the end of these weeks is calculated. Thus, the performance of the model in volatile market circumstances is tested. For these 25 weeks, the mean absolute yield error equals 9.90 basis points for the median firm. This illustrates that the latent factor model is capable of providing a good fit of credit spread data.

5. Estimation of the Default Jump Risk Premium

5.1 Estimation methodology

To set notation, define $Z_{j,t}$ as a variable that is equal to 1 if firm j defaults in the annual time interval $[t, t + 1]$ and 0 otherwise. Let $R_{j,t}$ denote the rating of firm j at time t . Estimation of the default jump risk premium μ is based on moment conditions for the conditional default rate, which is defined as the probability of default in year $t + n$, conditional upon no default between time t and $t + n$ and the rating at time t . In formulas, the moment conditions are

$$E^P(Z_{j,t+n} | R_{j,t} = R, Z_{j,t} + Z_{j,t+1} + \dots + Z_{j,t+n-1} = 0) = q_{n,R}(\mu, \phi),$$

$$R = \text{AA}, \text{A}, \text{BBB}, \quad n = 0, \dots, 14, \quad (10)$$

where $q_{n,R}(\mu, \phi)$ denotes the model-implied conditional default rate under the actual probability measure, and ϕ is a parameter vector that contains all other parameters of the factor model. This model-implied default rate can be calculated from the process for the actual default intensity $h_{j,t}^P$. Since $h_{j,t}^P = h_{j,t}^Q / \mu$, this default rate is a decreasing function of μ and a function of the factor model parameters ϕ . Appendix E derives the expression for $q_{n,R}(\mu, \phi)$.⁹

By confronting these model-implied conditional default rates with the observed default rates, μ can be estimated. Both Moody's and Standard &

⁹ A constant loss rate L is assumed. Since only the product $h_{j,t}^Q L$ enters the model, including a time-varying loss rate would most likely lower the estimated variation of $h_{j,t}^Q$ but hardly change its unconditional expectation. Given that the mean of $h_{j,t}^Q$ is the most important determinant of conditional default rates, the assumption of constant loss rates is not likely to have a large impact on the estimates for the default jump risk premium.

Poor's (S&P) provide average historical cumulative default rates per rating category, which are averages of cumulative default rates of cohorts of firms that are formed each year.¹⁰ Given that the data on credit spreads start in 1991 and end in 2000, one would ideally use default rates on the cohort that starts in 1991 up to the cohort starting in 2000. However, since defaults do not occur frequently, one needs a relatively long period to reliably estimate default probabilities. For example, in the 1991–2000 period, default rates are low relative to the 1970s and 1980s, so that using the 1991–2000 period would lead to very high but imprecise estimates for the risk premium μ . Moreover, if one would use the default rate data for the 1991–2000 period, cumulative default rates for more than a 10-year period cannot be used, while the long-maturity bonds provide information on these long-term default rates. Therefore, we use a longer data period to obtain historical default rates. This has the disadvantage that the credit spread data and the default rate data are not entirely based on the same sample period. From an econometric viewpoint the use of two partly overlapping data periods is not a problem, since this estimation is based on averages of conditional default rates and the assumption of stationarity.¹¹

This study uses both S&P and Moody's data. The S&P default rate data are based on cohorts starting in 1981 up to 2000 [Standard & Poor's Special Report (2001)]. For comparison, I also perform an estimation based on Moody's data, which are based on the 1970–2000 period [Moody's Special Comment (2001)]. In this case it is assumed that the Moody's and S&P ratings are the same. The cumulative default rates from 1 year up to 15 years are used. However, longer horizons are not used for two reasons. First, S&P does not provide default rate data for longer horizons, and, second, Table 1 shows that for the median firm the maximum maturity of the fitted bonds is 15.0 years. For both datasets the cumulative default rates are converted into yearly conditional default rates $q_{n,R}^{\text{Data}}$ (Appendix E discusses this transformation). These empirical conditional default rates provide consistent estimates of the conditional expectation in the left-hand side of Equation (10), given that these rates are fully characterized by the initial rating and the conditioning period n . Next, the first step of the generalized method of moments (GMM) [Hansen (1982)] is applied to Equation (10), and the sum of squared differences between

¹⁰ Each cohort corresponds to a given initial date and contains all firms with the same credit rating at this date. The firms in a cohort are followed from this date onwards, and default rates can be calculated for each cohort. The average default rates across cohorts are used to average out potential time or cohort effects.

¹¹ For the firms in this sample, credit ratings are used at the end of the sample period (February 2000). This may lead to conservative estimates for μ if the credit ratings of firms have deteriorated over the sample period. S&P data [Standard & Poor's Special Report (2001)] show that for 1991–2000 the total number of downgrades was only slightly larger than the number of upgrades, so that there is not necessarily a systematic rating drift over time.

the model-implied and observed conditional default rates is minimized

$$\min_{\mu} \left[\sum_{R=AA,A,B} \sum_{n=0}^{14} (q_{n,R}(\mu, \hat{\phi}) - q_{n,R}^{\text{Data}})^2 \right] \quad (11)$$

over μ , inserting the estimates for the factor model $\hat{\phi}$. Appendix E describes the calculation of the standard error of the estimate for μ , which incorporates both the estimation error in the parameters of the factor model and the estimation error in the historical default rates.

5.2 Estimation results

This subsection presents the results for the three models. Besides the model in Equations (1)–(6) that has been discussed so far (the “benchmark model”), an “upper bound” model and a “lower bound” model are also estimated. The upper bound model is obtained by excluding the tax and liquidity corrections from the benchmark model. This model is included to see how much of corporate bond excess returns can be explained using default risk factors only [as in Huang and Huang (2002)]. Given the complete exclusion of tax and liquidity effects, the estimate for μ can be regarded as an upper bound for the true value. Estimation of this model is performed as before, leaving out the tax correction and step two of the estimation procedure described in Section 4.1 (the liquidity correction). For brevity, the results for the intermediate estimation steps are not presented, but qualitatively the parameter estimates are similar to the estimates reported in Tables 4 and 5.

For the lower bound model, the tax and liquidity effects are again excluded from the benchmark model. However, the term structure of AA-rated firms is used, instead of the government bond term structure, to calculate credit spreads for A-rated and BBB-rated firms. Thus, the difference between the AA term structure and the government bond term structure is fully attributed to time-varying liquidity and tax effects. Since in reality AA-rated firms contain a small amount of credit risk, the estimate for the default event risk premium may be interpreted as a lower bound, as long as there is no systematic liquidity difference between AA-rated bonds and A/BBB-rated bonds. This approach is also used by Jarrow et al. (2001). Given that the model for risk-neutral intensities fully specifies the behavior of government yields and corporate bond yields for all relevant rating categories, it is only necessary to re-estimate the default event risk premium for this model.

We start with the results for the benchmark model. Figure 3 shows the influence of the risk premium μ on default probabilities. This figure plots yearly conditional default probabilities for the BBB rating category. The line “Risk-neutral: BM” depicts the risk-neutral default probabilities




Figure 3
Conditional default probabilities

The graph contains yearly conditional default probabilities for BBB-rated firms. The yearly conditional default probability is defined as the probability of default in the next year, given that no default has occurred before. The line “S&P Data” gives the historically estimated default probabilities, obtained from S&P data. The other lines all are model-implied default probabilities:

- “Risk-neutral: BM” refers to the risk-neutral probabilities for the benchmark model (BM).
- “ $\mu = 1$: BM” refers to the actual probabilities in case $\mu = 1$ for the BM.
- “ $\mu = 2.31$: BM” refers to the actual probabilities at the estimated value for μ (2.31) for the BM.

Furthermore, the actual probabilities for the upper bound and lower bound models (UBM and LBM) are given, both for the case where $\mu = 1$ and at the estimated values (5.51 and 1.80, respectively).

under the risk-neutral probability measure for the benchmark model. These are calculated by setting $\mu = 1$ and using the risk-neutral measure Q instead of the actual probability measure P for calculation of the default probabilities. The line “ $\mu = 1$: BM” shows what the model implies for actual default probabilities, assuming that there is no risk premium on the default jump ($\mu = 1$). The difference between these actual and risk-neutral probabilities is completely caused by the risk premia on the factors that drive the intensities, since the risk premia on these factors imply that the expectation of the path of the intensity under Q differs from the expectation under P .

Figure 3 also contains the empirical yearly conditional default rates based on S&P data from 1981–2000. These default rates are all below the default probabilities implied by the model with $\mu = 1$ (this is also the case for AA and A ratings). Thus, the risk premia on the intensity factors and the tax and liquidity effects cannot fully explain the level of observed default probabilities.

Next, the risk premium μ is estimated. Table 6a shows that this gives an estimate of 2.31 in case of the S&P data and 2.15 based on the Moody’s

Table 6
Estimates of the default event risk premium

	Benchmark model	Lower bound model	Upper bound model
(a) Estimates for the event risk premium μ			
S&P data: 1981–2000	2.31 (1.11)	1.80 (1.05)	5.51 (1.53)
Moody's data: 1970–2000	2.15 (1.03)	1.60 (0.97)	5.30 (1.44)
(b) Estimates for the event risk premium μ per rating category for S&P data: 1981–2000			
AA Category	1.83 (1.92)	—	7.18 (2.77)
A Category	2.61 (1.39)	2.03 (1.22)	5.32 (2.29)
BBB Category	2.37 (1.47)	1.67 (1.18)	4.78 (1.83)
(c) Benchmark model: Estimates for the event risk premium μ and tax rate			
	Risk premium μ	Tax rate (%)	
S&P data: 1981–2000	2.67	3%	
Moody's data: 1970–2000	2.46	3%	

The table contains estimates and standard errors for the default event risk premium, obtained using S&P data and Moody's data. (a) Estimates based on all rating categories together. (b) Estimates per rating category. The tax rate in the benchmark model in (a) and (b) is 4.875%. (c) Results for the joint estimation of the tax rate and the event risk premium. Estimation is performed as described in Section 5. Calculation of standard errors is described in Appendix E.

data.¹² Investors thus multiply the actual default probability with a factor of more than 2 for the pricing of corporate bonds. The standard errors are however quite large (1.11 and 1.03, respectively), due to the estimation error in the model parameters and estimation error in the historical default rates. Therefore, it is impossible to conclude with certainty that μ is larger than 1. This implies that once the intensity risk premia, and tax and liquidity effects, are corrected for, it is statistically not clear that the observed credit spreads are “too high” or that the average excess corporate bond returns are “too high.”

Although the statistical precision for $\hat{\mu}$ is not very large, including a default jump risk premium clearly improves the fit on historical default rates, as shown by the line “ $\mu = 2.31$: BM” in Figure 3.¹³ The economic contribution of all determinants becomes particularly clear from the plot in Figure 4, which gives the decomposition of expected excess corporate bond returns proposed in Equation (9). The graph shows that all components contribute to the expected return, but the individual contributions are quite different. For example, for a 10-year coupon-paying BBB bond, the default jump risk premium implies an excess return of 31 basis points

¹² The default jump risk premium has also been estimated using Duffee's model with only firm-specific factors to model intensities. In this case, the estimate for the default jump risk premium is of similar size as reported in Table 6. The advantage of the common factor model over Duffee's model is that it provides a decomposition of the total risk premium on the risk of intensity changes into common and firm-specific risk.

¹³ Results for AA and A ratings are similar. For all rating categories the observed 1-year, 2-year, and 3-year conditional default rates are slightly lower than the model-implied probabilities.




Figure 4

Corporate bond expected excess returns: benchmark model

For the median BBB-rated firm, the graph provides the model-implied estimate for the annualized expected return on a coupon-paying corporate bond with coupon of 7.6% (the average coupon for the median firm), for different maturities. The return is calculated in excess of the return on a government bond with the same maturity and coupon. Using the fact that a coupon bond is a portfolio of discount bonds, the conditional expected bond return for each week is calculated by applying Equations (9) and (A.1), inserting the factor values for that week, and subsequently averaging over all weeks. The average return is decomposed into four risk premia, as discussed in Section 2.2, and tax and liquidity effects are included according to the benchmark model. The tax effect is calculated by comparing the tax-corrected bond price [Equation (A.1)] with the uncorrected bond price and annualizing the percentage difference (given the bond maturity).

per year, while the interest rate and the common and firm-specific factor risk premia generate, respectively, -9 , 23 , and 4 basis points. Since credit spreads are negatively related to default-free rates, they provide a partial hedge against interest rate risk, which lowers the expected corporate bond return. The contribution of firm-specific and default-free factors is thus rather small. For this 10-year bond, the tax effect explains 33 basis points, and the liquidity factor almost 13 basis points. For higher rating categories, the tax and liquidity factors become more important since they (hardly) vary across rating categories. For example, for a 10-year AA-rated bond, the default jump risk premium generates an excess return of about 6 basis points, while tax and liquidity effects are essentially the same as for the BBB-rated bond. For a 10-year A-rated bond the jump premium is 12 basis points. These results can be compared with the informal calibration of Collin-Dufresne et al. (2003). They effectively set $\mu = 2$, which is close to the estimates in this article. Because they do not use historical default rates to estimate the default probability, but instead use a lower probability, their calibration results in a risk premium of 5 basis points for investment grade firms.

The common factors generate economically and statistically significant risk premia of about 9 (AA), 14 (A), and 23 (BBB) basis points for a 10-year coupon bond. In Section 1 it was shown that the common factors may capture similar effects as the contagion jumps of Collin-Dufresne et al. (2003). Collin-Dufresne et al. reach an upper bound for the contagion risk premium of 39 basis points, which compares well with the common factor risk premia in this article.

The results from this study can also be compared with those of Jarrow et al. (2001). They use the estimates of Duffee's (1999) model to compare model-implied default probabilities, with $\mu = 1$, with default probabilities that are implied by a Markov rating migration model. The migration model uses a one-year Moody's migration probability matrix. It turns out that the historically estimated cumulative default rates, as used in this article, are much lower than the cumulative default rates implied by this Markov model for rating migrations. This is evidence against the assumption of the Markov migration model that rating migrations are independent over time. It also implies that using the Markov migration model leads to a downward bias in the estimate for μ . Furthermore, Jarrow et al. find that the shape of the conditional default probabilities, as implied by the intensity model, is much flatter than the shape of the conditional default probabilities implied by the Markov migration model. Figure 3 shows that, if one uses the conditional default rates that are directly observed in the data, these shape differences disappear to a large extent.

To put the benchmark results into perspective, μ is estimated for the lower bound and upper bound model, which gives estimates of, respectively, 1.80 and 5.51 (Table 6a). The results for the upper bound model show that excluding tax and liquidity components has a dramatic effect on the estimate for μ . The estimate is very large in this case and the fit of observed default rates is less good (Figure 3). For the lower bound model, the estimate for μ is relatively close to the benchmark estimate. Given that default risk is small for AA-rated firms, so that a large part of the AA spread is likely to be caused by tax and liquidity effects, this small difference between benchmark and lower bound estimates supports the correction for tax and liquidity effects in the benchmark model.

Finally, a number of robustness checks are performed. First, μ is estimated for each rating category separately. Since μ is not expected to vary largely across rating categories, this can be interpreted as a cross-sectional test on the model. The results in Table 6b provide some support for the framework. The estimate is larger than 1 for all models and rating categories, and the variation across rating categories is not very large. Only for the upper bound model, the estimates for μ are significantly different from 1.

Second, the tax effect is analyzed in more detail. So far the results are based on a tax rate of 4.875%, as proposed by Elton et al. (2001). To check

whether this value is also appropriate for this sample, a grid search is performed to see which value for the tax rate gives the lowest value of the goal function in Equation (11). For each tax rate, μ is re-estimated. A joint estimation is infeasible since the factor model has to be re-estimated for each tax rate. Using 1% intervals, it was found that the optimal rate is 3% for both the S&P and Moody's data. The fact that this estimate is smaller than 4.875% can be understood as follows. Given that credit spreads are lowest for high-rated firms, the relative impact of the tax rate on model-implied default probabilities is highest for these high ratings. Table 6b shows that the AA rating has the lowest estimate for the jump risk premium. Therefore, decreasing the tax rate and increasing μ leads to a better fit of default probabilities across ratings. Unreported results show that, at the tax rate of 3%, the difference between the estimates for μ across rating categories is indeed smaller.

6. Concluding Remarks

The main contribution of this article is that it provides a careful empirical decomposition of corporate bond returns into several factors, including a risk premium associated with the jump in prices in case of a default event. This risk premium explains an economically significant part of BBB corporate bond returns, but the statistical evidence is inconclusive about the existence of a jump risk premium. Liquidity and tax effects, and a risk premium on market-wide credit spread movements are other important determinants of the expected returns of investment grade bonds. The article thus provides a first indication that default jump risk may not be fully diversifiable. However, more research is needed to obtain precise estimates of the jump risk premium. This may be possible by combining information on risk-neutral intensities in credit spread data with (model-based) estimates of actual default probabilities on the firm level [Das et al. (2002)].

There are several extensions worth exploring in future research. It would be interesting to see whether utility-based models can explain the size of the estimated jump risk premium [see El Karoui and Martellini (2001)]. Another extension would be to include high-yield bonds in the analysis. Finally, Keswani (1999) and Duffie et al. (2003) study the pricing of defaultable sovereign debt. The model in this article can be used to analyze the joint pricing of sovereign debt spreads of many countries.

Appendix A: Tax Correction on Corporate Bond Prices

This appendix estimates the tax correction in the intensity-based framework. Let t_s denote the state tax rate and t_g the federal tax rate. As shown by Elton et al. (2001), the effective tax rate is $t_s(1 - t_g)$, which is denoted by τ . This tax rate changes the corporate bond price in two ways. First, if the coupon size is denoted by C , the net coupon received is $(1 - \tau)C$. Second, there is a

tax recovery on the default loss if the firm defaults, which changes the loss rate to $(1 - \tau)L$. It is assumed that this loss rate applies to the net coupon payments. This leads to the following valuation equation for a corporate bond that has n coupon payments and associated payment dates $T_1, \dots, T_n$

$$V_j(t, T, n) = (1 - \tau) C \sum_{i=1}^n E_t^Q \left[\exp \left(- \int_t^{T_i} (r_s + \beta_{0j}^l + \beta_{1j}^l I_s + h_{j,s}^Q (1 - \tau)L) ds \right) \right] + E_t^Q \left[\exp \left(- \int_t^{T_n} (r_s + \beta_{0j}^l + \beta_{1j}^l I_s + h_{j,s}^Q (1 - \tau)L) ds \right) \right]. \quad (\text{A.1})$$

Appendix B: Parameter Identification

From Dai and Singleton (2000) and Duffee (2002) it follows that all parameters related to the default-free and firm-specific factors can be identified. This appendix analyzes which parameters in the common factor processes can be identified. In Equation (5) it is shown that the contribution of the common factor i to the instantaneous spread of firm j is given by $\bar{F}_{ij,t} = \gamma_{ij} F_{i,t}$.

The process of $\bar{F}_{ij,t}$ under the risk-neutral measure Q is given by

$$d\bar{F}_{ij,t} = (\gamma_{ij} \kappa_i^F \theta_i^F - (\kappa_i^F + \lambda_i^F) \bar{F}_{ij,t}) dt + \sqrt{\gamma_{ij}} \sigma_i^F \sqrt{\bar{F}_{ij,t}} d\hat{W}_{i,t}^F, \quad i = 1, \dots, K. \quad (\text{B.1})$$

The process under the true probability measure is obtained by removing λ_i^F from Equation (B.1) and replacing the Q -Brownian motion $\hat{W}_{i,t}^F$ with a P -Brownian motion $\hat{W}_{i,t}^P$. The identifiable parameters are the parameters in the processes under P and Q , $(\gamma_{ij} \kappa_i^F \theta_i^F, \kappa_i^F, \kappa_i^F + \lambda_i^F, \sqrt{\gamma_{ij}} \sigma_i^F)$. These four reduced-form parameters are a function of five structural parameters. For another firm k the identifiable parameters are $(\gamma_{ij} \kappa_i^F \theta_i^F, \kappa_i^F, \kappa_i^F + \lambda_i^F, \sqrt{\gamma_{ij}} \sigma_i^F)$. It then follows that, besides κ_i^F and λ_i^F , it is not possible to recover the remaining structural parameters from the reduced form parameter estimates, and that normalizing θ_i^F , $i = 1, \dots, K$, solves this identification problem. θ_i^F is normalized to 50 basis points.

Appendix C: Expected Corporate Bond Returns

This appendix derives the expression for the instantaneous expected excess return on a corporate bond. The derivation is based on Yu (2002). First, for a moment, the possibility of default is neglected. As discussed in Section 2, the model implies that corporate bond prices are an exponential-affine function of the underlying factors. Applying Ito's lemma to this equation gives the expected return in case of no default

$$E_t \left[\frac{dV_j(t, T)}{V_j(t, T)} \middle| \text{no default} \right] = \left[r_t + s_{j,t} - \sum_{i=1}^K (T - t) B_{ij}(T - t) \lambda_i^F F_{i,t} - (T - t) C_j(T - t) \lambda_j^G G_{j,t} + \tilde{A}(\alpha_j + \delta, X_t, T - t) - (T - t) D_j(T - t) \lambda_j^l l_t \right] dt. \quad (\text{C.1})$$

Of course, Equation (C.1) neglects the loss in case of default. In a small time interval $(t, t + \Delta t)$ the default probability approximately equals $h_{j,t}^P \Delta t$ and the loss in case of default equals $L \cdot V_j(t, T)$. Then, the expected return over the next time interval approximately equals

$$\left\{ (1 - h_{j,t}^P \Delta t) E_t \left[\frac{\Delta V_j(t, T)}{V_j(t, T)} \middle| \text{no default} \right] + (h_{j,t}^P \Delta t) (-L) \right\} / \Delta t. \quad (\text{C.2})$$

This expression becomes exact if $\Delta t \rightarrow 0$, which gives the instantaneous expected return

$$r_t + s_{j,t} - \sum_{i=1}^K (T-t) B_{ij} (T-t) \lambda_i^F F_{i,t} - (T-t) C_j (T-t) \lambda_j^G G_{j,t} + \tilde{A}(\alpha_j + \delta_j X_t, T-t) - (T-t) D_j (T-t) \lambda_j^L l_t - h_{j,t}^P L. \quad (C.3)$$

The return in excess of a default-free bond (with the same maturity) is obtained by subtracting the expected return in Equation (7) from the expression in Equation (C.3), which gives Equation (9).

Appendix D: The Kalman Filter Setup

This appendix briefly describes the setup for the extended Kalman filter estimation of the affine pricing models for default-free and defaultable bonds, using notation that is unrelated to the notation in the main text. Let F_t be a K -dimensional vector with process under P given by

$$dF_t = \Lambda(\theta - F_t)dt + \Sigma(\alpha + B'F_t)_d^{1/2} dW_t. \quad (D.1)$$

Here Λ , Σ , and B are K by K matrices, θ and α are K -dimensional vectors, and W_t is a K -dimensional vector of independent Brownian motions. The term $(\alpha + B'F_t)_d^{1/2}$ is a diagonal matrix with diagonal elements given by the square root of the elements of the vector $(\alpha + B'F_t)$. Without loss of generality we take the matrix Λ to be diagonal, $\Lambda = \text{diag}(\kappa_1, \dots, \kappa_K)$, since one can transform any affine model into a model with diagonal Λ by rotating the factors.

Discretizing (D.1) gives the following transition equation

$$F_{t+h} = E_t[F_{t+h}|F_t] + \eta_{t+h}, \quad V_t(\eta_{t+h}) = V(F_{t+h}|F_t). \quad (D.2)$$

Refer to De Jong (2000) for expressions of the conditional expectation and variance in (D.2).

This dataset consists of yields on coupon bonds. Let Y_t denote an N -dimensional vector with the observed yields at time t . The measurement equation for Y_{t+h} is then given by

$$Y_{t+h} = z(F_{t+h}) + \varepsilon_{t+h}, \quad V_t(\varepsilon_{t+h}) = H_t. \quad (D.3)$$

Here $z(F_{t+h})$ is a function that relates the yields to the factors and ε_{t+h} is a zero-expectation measurement error that is uncorrelated with F_{t+h} . Since coupon-yields are used, the function $z(F_{t+h})$ is nonlinear. As in Duffee (1999), a Taylor approximation of this function around the one-period forecast of F_{t+h} is used to linearize the model. Using Equations (D.2) and (D.3) the likelihood function can be constructed [see De Jong (2000)].

For the default-free model, it is assumed that the covariance matrix of ε_t is a diagonal matrix, $H_t = \text{diag}(\sigma_{\varepsilon,1}^2, \dots, \sigma_{\varepsilon,N}^2)$. Thus, the common factor model for the intensity process is estimated twice. In the first estimation, it is assumed that $H_t = \sigma_{\varepsilon}^2 I_N$. After this first estimation the firm-specific factor parameters are estimated using the residuals from the common factor model. For each firm, there is a state-space model for these residuals again with a diagonal measurement error covariance matrix. Given the estimation results for the firm-specific factors, the common factor model is now re-estimated using a block-diagonal structure for H_t , where each block relates to all corporate bond yields of a single firm. Yield measurement errors across firms are assumed to be uncorrelated. For each block the conditional covariance matrix of yield residuals $V_t(\varepsilon_{t+h})$ is used, as implied by the estimates for the firm-specific factor process of this firm.

An important issue for the estimation of the common factor model is the presence of missing observations. Each week, always some corporate yields are observed in the dataset. Therefore, to construct the weekly likelihood contribution, all available observations for that week are used. For the estimation of the firm-specific factors the treatment of missing

observations is different because there are weeks with no observations at all for a single firm. As in Duffee (1999), the length of the time interval h is allowed to vary over time in this case.

Assuming that all factors follow stationary processes under the true probability measure, the unconditional expectation and (co)variances of the factors can be used to initiate the Kalman filter. Refer to De Jong (2000) for all equations in the Kalman filter recursion.

Finally, it can be noted that this estimation method is quasi maximum likelihood because the conditional distribution of F_{t+h} conditional on F_t is not normal. Moreover, the conditional variance in this distribution depends on the unknown values F_t , which makes the QML estimator based on the Kalman filter, strictly speaking, inconsistent. Simulation experiments by Duan and Simonato (1999) and De Jong (2000) show that the induced biases are small. Consistent parameter estimates can be obtained by using the efficient method of moments [EMM, Gallant and Tauchen (1996)], combined with the semi-nonparametric (SNP) method of Gallant and Tauchen (1992). Duffee and Stanton (2000) compare EMM/SNP estimation of affine term structure models with the Kalman filter QML estimation. They document considerable small-sample biases for the EMM/SNP method, and conclude that “for reasonable sample sizes, the results strongly support the choice of the Kalman filter.”

Appendix E: Estimation of the Default Jump Risk Premium

This appendix first describes the calculation of the model-implied conditional default rate $q_{n,R}(\mu, \phi)$. Using the following result for the actual probability $p_{n,j}(t, \mu, \phi)$ that a firm defaults within the next n years, conditional upon that no default has occurred yet

$$\begin{aligned} p_{n,j}(t, \mu, \phi) &\equiv E[Z_{j,t} + Z_{j,t+1} + \dots + Z_{j,t+n-1} | R_j = R, \{F_{1,t}, F_{2,t}, G_{j,t}\}] \\ &= 1 - E_t^P \left[\exp \left(- \int_t^{t+n} h_{j,s}^P ds \right) \right] = 1 - E_t^P \left[\exp \left(- \int_t^{t+n} \frac{h_{j,s}^Q}{\mu} ds \right) \right]. \end{aligned} \quad (E.1)$$

The expectation in Equation (E.1) can be calculated explicitly because the process for $h_{j,t}^Q$ is affine. Note that in Equation (E.1) the time-dependence of the firm's rating has been taken out, since the model does not allow for rating changes. The firm's rating matters since the common factor parameters vary across rating categories. To average out the factor values, Equation (E.1) is averaged over all factor values (obtained from the Kalman filter) for the 1991–2000 period. The resulting probabilities are denoted by $p_{n,j}(\mu, \phi)$. Instead of this average, the unconditional expectation of Equation (E.1) can also be used. Since the average fitted factor values are mostly lower than the unconditional means, this will most likely lead to higher estimates for μ . Finally, the yearly conditional default rates $q_{n,j}(\mu, \phi) \equiv 1 - (1 - p_{n+1,j}(\mu, \phi)) / (1 - p_{n,j}(\mu, \phi))$ are calculated, and these conditional default probabilities are averaged over all firms in a given rating category to obtain $q_{n,R}(\mu, \phi)$.

The standard error for the estimate for μ is calculated by incorporating both the estimation error in the historically estimated default rates and the estimation error in the factor model parameter estimates.

To correct the asymptotic variance of the estimator for μ for the estimation error in ϕ , $q(\mu, \phi)$ is defined as $(q_{0,AA}(\mu, \phi), \dots, q_{14,AA}(\mu, \phi), q_{0,A}(\mu, \phi), \dots, q_{14,A}(\mu, \phi), q_{0,BBB}(\mu, \phi), \dots, q_{14,BBB}(\mu, \phi))'$ and, similarly, the sample counterpart vector q^{Data} as well as $\Gamma = \partial q(\mu, \phi) / \partial \mu$ are also defined. The estimation method in Equation (11) is the first step of GMM [Hansen (1982)]. The GMM first-order conditions imply that [see, e.g., Gouriéroux and Monfort (1995)]

$$\Gamma'(q^{\text{Data}} - q(\hat{\mu}, \hat{\phi})) \xrightarrow{p} 0. \quad (E.2)$$

Using a first-order Taylor expansion for $q(\mu, \phi)$ (at the parameter estimates) and rewriting gives

$$(\hat{\mu} - \mu) \xrightarrow{p} \frac{\Gamma'}{\Gamma'\Gamma} (q^{\text{Data}} - q(\mu, \phi)) - \frac{\Gamma'}{\Gamma'\Gamma} \frac{\partial q(\mu, \phi)}{\partial \phi'} (\hat{\phi} - \phi). \quad (\text{E.3})$$

Given the appropriate regulatory conditions, both terms on the right-hand side have asymptotically a normal distribution (if scaled with the square root of the number of observations). It is assumed that the estimation error in the historical default rates is uncorrelated with the estimation error in the parameters of the factor model. For both terms on the right-hand side of Equation (E.3), the calculation of the asymptotic covariance matrix is explained below. The standard error for $\hat{\mu}$ can then be calculated using Equation (E.3).

First, the variances of the observed conditional default rates $q_{n,R}^{\text{Data}}$ are calculated, by assuming that in each year defaults are independently generated by a binomial distribution with probability $q_{n,R}^{\text{Data}}$. S&P and Moody's report the number of firms per cohort (and per rating category), and these are summed over all relevant cohorts. For example, for the default probability in the first year, $q_{0,R}^{\text{Data}}$, the sum is taken over all cohorts, while for the 14-year conditional default probability $q_{14,R}^{\text{Data}}$, the sum is taken only over the cohorts starting in 1981 up to 1986 in case of S&P data, and over the cohorts starting in 1970 up to 1986 in case of Moody's data. The variance of this default probability estimate can then be determined as $q_{n,R}^{\text{Data}}(1 - q_{n,R}^{\text{Data}})/N_R$, where N_R is the total number of firms in all yearly cohorts of a given rating category.

A more difficult issue is the correlation between the observed default rates. The issue is that a large fraction of the firms in the AA-cohort that starts in, say, 1982, will also be present in the AA-cohort that starts in 1983. To be conservative, it is assumed that there is perfect correlation between the estimated default rates for different n within each rating category, and that there is zero correlation between the estimated default rates across rating categories. Given these assumptions, it is possible to calculate the full covariance matrix of the default rates $q_{n,R}^{\text{Data}}$.

Second, the covariance matrix of the second part of the right-hand side of Equation (E.3) is calculated. It is assumed that the four sets of parameters (default-free, liquidity factor, common factor, and firm-specific factor parameters) are independent from each other. Thus, it is possible to construct the block-diagonal covariance matrix of all these estimates. Using the delta-method, one can construct the covariance matrix of the model-implied conditional default rates for all firms, and finally the covariance matrix of the average conditional default rates of all firms in a given rating category.

References

- Bierens, H., J. Huang, and W. Kong, 2003, "An Econometric Model of Credit Spreads with Re-balance, ARCH and Jump Effects," working paper, Penn State University.
- Chen, L., and J. Wei, 2002, "An Indirect Estimate of Transaction Costs for Corporate Bonds," working paper, Michigan State University.
- Collin-Dufresne, P., R. Goldstein, and J. Helwege, 2003, "Is Credit Event Risk Priced? Modelling Contagion via the Updating of Beliefs," working paper, Carnegie Mellon University.
- Collin-Dufresne, P., R. S. Goldstein, and J. S. Martin, 2001, "The Determinants of Credit Spread Changes," *Journal of Finance*, 56, 2177–2207.
- Chakravarty, S., and A. Sarkar, 1999, "Liquidity in US Fixed Income Markets: A Comparison of the Bid-Ask Spread in Corporate, Government and Municipal Bond Markets," working paper, Purdue University and Federal Reserve Bank New York.
- Dai, Q., and K. J. Singleton, 2000, "Specification Analysis of Affine Term Structure Models," *Journal of Finance*, 55, 1943–1978.

- Das, S., L. Freed, G. Geng, and N. Kapadia, 2002, "Correlated Default Risk," working paper, Santa Clara University.
- De Jong, F., 2000, "Time-Series and Cross-section Information in Affine Term Structure Models," *Journal of Economics and Business Statistics*, 18, 300–318.
- Delianedis, G., and R. Geske, 2001, "The Components of Corporate Credit Spreads: Default, Recovery, Tax, Jumps, Liquidity, and Market Factors," working paper, UCLA.
- Duan, J., and J. G. Simonato, 1999, "Estimating Exponential-Affine Term Structure Models by Kalman Filter," *Review of Quantitative Finance and Accounting* 13, 11–135.
- Duffee, G., 1999, "Estimating the Price of Default Risk," *Review of Financial Studies*, 12, 197–226.
- Duffee, G., 2002, "Term Premia and Interest Rate Forecasts in Affine Models," *Journal of Finance*, 57, 405–443.
- Duffee, G., and R. Stanton, 2000, "Estimation of Dynamic Term Structure Models," working paper, Haas School of Business, University of California.
- Duffie, D., and R. Kan, 1996, "A Yield-Factor Model of Interest Rates," *Mathematical Finance*, 6, 379–406.
- Duffie, D., L. H. Pedersen, and K. J. Singleton, 2003, "Modeling Sovereign Yield Spreads: A Case Study of Russian Debt," *Journal of Finance*, 58, 119–159.
- Duffie, D., and K. J. Singleton, 1999, "Modeling Term Structures of Defaultable Bonds," *Review of Financial Studies*, 12, 687–720.
- El Karoui, N., and L. Martellini, 2001, "A Theoretical Inspection of the Market Price for Default Risk," working paper, University of Southern California.
- Elton E. J., M. J. Gruber, D. Agrawal and C. Mann, 2001, "Explaining the Rate Spread on Corporate Bonds," *Journal of Finance*, 56, 247–277.
- Elton E. J., M. J. Gruber, D. Agrawal and C. Mann, 2002, "Factors Affecting the Valuation of Corporate Bonds," working paper, New York University.
- Gallant, A. R., and G. E. Tauchen, 1992, "A Nonparametric Approach to Nonlinear Time Series Analysis: Estimation and Simulation," in D. Brillinger, P. Caines, J. Geweke, E. Parzen, M. Rosenblatt, and M. S. Taquq (eds.) *New Directions in Time Series Analysis, Part II*, pp. 71–92 Springer-Verlag, New York.
- Gallant, A. R., and G. E. Tauchen, 1996, "Which Moments to Match?," *Econometric Theory*, 12, 657–681.
- Gourieroux, C., and A. Monfort, 1995, *Statistics and Econometric Models*, Cambridge University Press, Cambridge.
- Hansen, L.P., 1982, "Large Sample Properties of Generalized Methods of Moments Estimators," *Econometrica*, 50, 1029–1054.
- Houweling, P., A. Mentink, and T. Vorst, 2003, "How to Measure Corporate Bond Liquidity?," working paper, Erasmus University Rotterdam.
- Huang, J., and M. Huang, 2002, "How Much of the Corporate-Treasury Yield Spread is Due to Credit Risk?," working paper, Penn State University.
- Janosi, T., R. Jarrow, and Y. Yildirim, 2002, "Estimating Expected Losses and Liquidity Discounts Implicit in Debt Prices," *The Journal of Risk*, 5(1).
- Jarrow, R. A., D. Lando, and F. Yu, 2001, "Default Risk and Diversification: Theory and Applications," working paper. Cornell University.
- Jarrow, R. A., and S. M. Turnbull, 1995, "Pricing Derivatives on Financial Securities Subject to Credit Risk," *Journal of Finance*, 50, 53–86.

- Keswani, A., 1999, "Estimating a Risky Term Structure of Brady Bonds," working paper, London Business School.
- Leland, H., 2002, "Predictions of Expected Default Frequencies in Structural Models of Debt," working paper, UCLA.
- Liu, J., F. Longstaff, and R. Mandell, 2001, "The Market Price of Credit Risk: An Empirical Analysis of Interest Rate Swap Spreads," working paper, UCLA.
- Longstaff, F. A., and E. S. Schwartz, 1995, "A Simple Approach to Valuing Risky and Floating Rate Debt," *Journal of Finance*, 50, 789–819.
- Madan, D. B., and H. Unal, 1998, "Pricing the Risks of Default," *Review of Derivatives Research*, 2, 121–160.
- Moody's Special Comment, 2001, "Default and Recovery Rates of Corporate Bond Issuers: 2000," Moody's Investors Service, Global Credit Research.
- Newey, W. K., and K. D. West, 1987, "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.
- Pearson, N. D., and T. S. Sun, 1994, "Exploiting the Conditional Density in Estimating the Term Structure: An Application of the CIR model," *Journal of Finance*, 49, 1279–1304.
- Standard & Poor's Special Report, 2001, "Ratings Performance 2000," Standard & Poor's Risk Solutions.
- Yu, F., 2002, "Decomposing the Expected Return on Corporate Bonds," *Journal of Fixed Income*, 12, 69–81.
- White, H., 1982, "Maximum Likelihood Estimation of Misspecified Models," *Econometrica*, 50, 1–25.

Copyright of Review of Financial Studies is the property of Society for Financial Studies and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.