Modeling Expected Return on Defaultable Bonds

FAN YU

FAN YU
is an assistant professor of finance at the Graduate School of Management at the University of California at Irvine.
fanyu@uci.edu

Research in the past few years has greatly enhanced our understanding of default risk. For example, the popular reduced-form intensity approach, following Jarrow and Turnbull [1995], Lando [1998], and Duffie and Singleton [1999], among others, has been used to model corporate-Treasury spreads (Duffee [1999]), LIBOR-swap spreads (Collin-Dufresne and Solnik [2001]), swap-Treasury spreads (Duffie and Singleton [1997], Liu, Longstaff, and Mandell [2000]), and sovereign yield spreads (Duffie, Pedersen, and Singleton [2001]). These studies focus on fitting and interpreting the time series properties of the spreads, and on that measure they are quite successful. Others such as Bakshi, Madan, and Zhang [2001] have undertaken to examine the out-of-sample pricing and hedging performance of such models, also with encouraging results.

Conspicuously missing in this literature is a discussion of the default risk premium, or, more generally, the expected return on defaultable bonds. In the default-free term structure literature, it is recently realized that the risk premium specification in term structure models should be flexible enough to produce time-varying expected returns consistent with empirically observed bond excess returns (Duffee [2002]). Although much is still unknown about the empirical features of excess returns on defaultable bonds (perhaps because of the difficulty in measuring returns through default and/or restructuring), a general intensity-based default model could yield valuable guidance for empirical research on this topic.

It may be tempting to think that the issue of defaultable bond expected return has been rendered moot by the work of Duffie and Singleton [1999], which shows that under certain technical conditions, defaultable bonds are priced pre-default with an adjusted short rate that can be modeled as an affine function of diffusion state variables just as in the default-free case.[1] Indeed, for valuation purposes the default-free case and the defaultable case are methodologically indistinguishable—the tractability of affine models, for example, can be exploited in both cases.

Unfortunately, this elegant result gives the (unintended) impression that specification of the default risk premium and defaultable bond expected return can also proceed as in default-free term structure models. The liberal use in the empirical literature of terms such as "price of default risk" and "credit premium," representing the drift adjustments on the affine diffusion state variables underlying the risk-neutral default intensity, is evidence of this confusion.

Jarrow, Lando, and Yu [2001] (JLY) consider specification of the default risk premium in reduced-form models. They show that there are generally two types of default risk premium: one due to systematic variations in the default intensity, and the other a jump risk premium on the default event. When default risk is what they characterize as "conditionally diversifiable," the default event risk premium will disappear, and one can specify the default risk premium using drift adjustments as in the default-free case. JLY propose using joint data on default rates and defaultable bond prices to estimate this default event risk premium.[2]

This article clarifies the implication of JLY's risk premium specification for modeling the expected return on defaultable bonds. Despite the strong analogy between default-free and defaultable bond pricing, the similarity between the two sets of securities stops when it comes to expected return. Intuitively, to compute the expected return on defaultable bonds, one has to know the product of the likelihood of default and how much of bond value is expected to be lost post-default, something appropriately termed the "mean loss rate" under the physical measure. In Duffie and Singleton [1999], however, pre-default bond prices are dependent on the mean loss rate under the risk-neutral measure. This subtle difference caused by the measure change is absent in the modeling of default-free expected returns.

Using a simple illustrative example from the affine class, we show that the expected return on a defaultable bond can generally be decomposed into three parts. The first component is the expected return on an otherwise identical default-free bond. The second component is related to the parameters of the risk-neutral mean loss rate and their risk adjustments under the physical measure. It can be understood as a compensation for variations in credit risk generated by the state variables.

The third and so far ignored component is the difference between the risk-neutral and the physical mean loss rates, and can be thought of as a compensation demanded by investors for bearing the risk of the (systematic or non-diversifiable) default event. Incidentally, this result can also be interpreted along the lines of "survival bias" frequently mentioned in studies of stock returns (Brown, Goetzmann, and Ross [1995], Li and Xu [2000]). If one uses only bonds not in default in the estimation of expected return, or if one computes expected return based on the pre-default pricing formula given by Duffie and Singleton [1999], the expected return will be overstated. The extent of the overstatement is related to the physical mean loss rate of the bond.

I. DECOMPOSITION OF EXPECTED RETURN

We first present a simple calculation of the expected return on defaultable bonds. To illustrate the conceptual distinction from default-free bonds, a standard affine model of the default-free term structure is first considered. Although our derivation relates to the Cox, Ingersoll, and Ross [1985] specification, the decomposition result can be broadly extended to the entire affine class.

Default-Free Case

The benchmark case for comparison purposes is a one-factor CIR model of the default-free term structure. In this model, the dynamics of the short rate r_t under the physical measure P and the risk-neutral measure Q are assumed to be:

$$dr_t = \kappa(\theta - r_t)\,dt + \sigma\sqrt{r_t}dW_t, \text{ under } P \tag{1}$$

$$dr_t = (\kappa+\nu)\left(\frac{\kappa\theta}{\kappa+\nu} - r_t\right)dt + \sigma\sqrt{r_t}d\widetilde{W}_t, \text{ under } Q \tag{2}$$

where κ, θ, and ν are constants, and W_t and $\widetilde{W}_t$ are P- and Q-Brownian motion, respectively.

In this case, it is well known that zero-coupon bond prices can be expressed as an exponentially affine function of the short rate:

$$p(t,T) = H_1 \exp(-H_2 r_t) \tag{3}$$

where

$$H_1 = \left(\frac{2\gamma e^{(\gamma+\kappa+\nu)(T-t)/2}}{(\gamma+\kappa+\nu)\left(e^{\gamma(T-t)}-1\right)+2\gamma}\right)^{2\kappa\theta/\sigma^2} \tag{4}$$

$$H_2 = \frac{2\left(e^{\gamma(T-t)}-1\right)}{(\gamma+\kappa+\nu)\left(e^{\gamma(T-t)}-1\right)+2\gamma} \tag{5}$$

$$\gamma = \sqrt{(\kappa+\nu)^2+2\sigma^2} \tag{6}$$

It is then straightforward to apply Ito's lemma to Equation (3), yielding the dynamics as follows for $p(t, T)$:

$$\frac{dp(t,T)}{p(t,T)} = (r_t + H_2\nu r_t)\,dt - H_2\sigma\sqrt{r_t}dW_t \tag{7}$$

Thus the instantaneous expected return is

$$r_t + H_2\nu r_t \tag{8}$$

which is an affine function of r_t.

This example shows that the expected return depends on both the risk-neutral dynamics of r_t (through H_2) and the parameter ν that determines the physical dynamics of r_t. For this reason ν is often referred to as the market price of interest rate risk.

Defaultable Case

Now consider computation of the expected return on defaultable bonds. For simplicity, the one-factor CIR model for the default-free term structure is maintained. Consideration for default risk is achieved by constructing a risk-neutral mean loss rate $\widetilde{s}_t$ that follows a square-root diffusion:

$$d\widetilde{s}_t = (\kappa' + \nu')\left(\frac{\kappa'\theta'}{\kappa' + \nu'} - \widetilde{s}_t\right)dt + \sigma'\sqrt{\widetilde{s}_t}d\widetilde{W}'_t \quad (9)$$

where κ', θ', and ν' are constants, and $\widetilde{W}'_t$ is a Brownian motion independent of $\widetilde{W}_t$ under Q. Similarly, it is assumed that the P-dynamics of $\widetilde{s}_t$ are given by

$$d\widetilde{s}_t = \kappa'(\theta' - \widetilde{s}_t)\,dt + \sigma'\sqrt{\widetilde{s}_t}dW'_t \quad (10)$$

where W'_t is a Brownian motion independent of W_t under P.

This specification of the mean loss rate ignores the empirical observation that yield spreads change systematically with the default-free term structure, but it suffices for this simple illustration.[3]

As shown by Duffie and Singleton [1999], under certain technical assumptions defaultable bonds can be valued like default-free bonds—one needs only to modify the short rate used for discounting from r_t to $R_t = r_t + \widetilde{s}_t$.[4] The risk-neutral mean loss rate $\widetilde{s}_t$ is equal to $\widetilde{\lambda}_t(1 - \delta_t)$ where $\widetilde{\lambda}_t$ is the risk-neutral default intensity, and δ_t is the recovery rate as a fraction of pre-default market value of the bond. In this setting, once we make the assumption that R_t is an affine function of diffusion state variables, defaultable bond pricing collapses to the default-free case.

Naive Derivation. In this example, it can be shown that the defaultable zero-coupon bond price is given by

$$v(t,T) = H_1 H'_1 \exp\left(-H_2 r_t - H'_2\widetilde{s}_t\right) \quad (11)$$

where H_1 and H_2 are defined by Equations (4)-(6), and H'_1 and H'_2 are the corresponding prime versions. Thinking of to the derivation of Equation (7), it is tempting to conclude that the instantaneous expected return on the defaultable bond is

$$r_t + \widetilde{s}_t + H_2\nu r_t + H'_2\nu'\widetilde{s}_t \quad (12)$$

which follows from the dynamics of $v(t, T)$:

$$\frac{dv(t,T)}{v(t,T)} = \left(r_t + \widetilde{s}_t + H_2\nu r_t + H'_2\nu'\widetilde{s}_t\right)dt - H_2\sigma\sqrt{r_t}dW_t - H'_2\sigma'\sqrt{\widetilde{s}_t}dW'_t \quad (13)$$

The derivation of expected return is similar to that of Liu, Longstaff, and Mandell [2000] (LLM), who use a multifactor Vasicek specification instead. It is not difficult to see that the expression in Equation (12) is incomplete, however. Take the simple example that the market price of risk is zero, i.e., $\nu = 0$ and $\nu' = 0$. In this case, the defaultable bond is priced under risk-neutrality, and the instantaneous expected return ought to be just the risk-free rate r_t. Instead, Equation (12) gives $r_t + \widetilde{s}_t$.

Hence, an important piece is missing from the derivation. We show that this missing piece is related to the expected return conditional on a default event.

Correct Derivation. The intuition behind why Equations (12)-(13) can go wrong is simple, once it is realized that although Duffie and Singleton's methodology can reduce the pricing of defaultable bonds to that of default-free bonds, it is applicable only before the default event. In deriving expected returns on defaultable bonds, however, one cannot avoid the question of what happens to bond value in default and how likely default will occur.

We present a somewhat heuristic derivation of the instantaneous expected return on a defaultable bond. To start, it is assumed that under the physical measure P, the default intensity is given by λ_t.[5] Over the interval $(t, t + \Delta t)$, the likelihood of default is given by $\lambda_t\Delta t$. Thus the return over the interval is:

$$\left.\frac{\Delta v(t,T)}{v(t,T)}\right|_{\text{no default}}(1 - \lambda_t\Delta t) + \frac{\delta_t v(t-,T) - v(t-,T)}{v(t-,T)}\lambda_t\Delta t \approx \left.\frac{\Delta v(t,T)}{v(t,T)}\right|_{\text{no default}} - (1 - \delta_t)\lambda_t\Delta t. \quad (14)$$

where the first term becomes Equation (13) in the limit as $\Delta t \to 0$.

Hence the instantaneous expected return is slightly modified from Equation (12) as

$$r_t + H_2\nu r_t + H_2^{'}\nu'\widetilde{s}_t + \widetilde{s}_t - s_t \tag{15}$$

where $s_t = \lambda_t(1-\delta_t)$ is the mean loss rate under the physical measure.

The difference between the two equations is essentially a consequence of survival bias—expected returns are bound to be overstated when they are estimated using only the surviving subsample. The extent of the survival bias is related to the physical default intensity λ_t and not the risk-neutral intensity $\widetilde{\lambda}_t$. While the issue of survival bias is largely an afterthought in the literature on stock returns, it is of paramount importance here, since our focus is precisely on the risk and return associated with default and bankruptcy.[6]

Interpretation of the Jarrow-Lando-Yu Default Risk Premium Specification. The instantaneous expected return on defaultable bonds, given in Equation (15), can be seen to consist of three components. The first component is the return on an otherwise identical default-free bond (the first two terms).[7] The second is an affine function of state variables (third term). The third is the difference between the risk-neutral and physical expected loss rate.

These components can help us understand the default risk premium specification in the recent work of JLY. They consider an intensity-based model of multiple issuers whose defaults are independently conditional on a set of state variables. By using a diversification argument similar to that of the original arbitrage pricing theory, they are able to show asymptotic equivalence between risk-neutral and physical default intensities.

The implication for Equation (15) is that the third component will disappear because $\lambda_t = \widetilde{\lambda}_t$. In such cases, investors do not demand any compensation for bearing the risk of the default event. Instead, they are compensated because of variation in default risk, which enters Equation (15) as the second component. It is in this sense that the parameter ν' can be considered a "default risk premium."

When the risk inherent in the actual default event cannot be diversified away, one has to consider the third component explicitly. This is made possible because of Girsanov-type results showing that under an equivalent change of measure, point process intensities are linked by a strictly positive predictable process μ_t through the relation $\widetilde{\lambda}_t = \mu_t\lambda_t$. Hence the compensation for the default event becomes $(\mu_t - 1)s_t$.

It is in this sense that the parameter μ_t is labeled a "default event risk premium." Since investors' risk aversion as to the default event implies a μ_t greater than one, ignoring the third component will lead to an underestimation of the expected return on defaultable bonds.

Non-Default Sources of Expected Return. Duffie and Singleton [1999] show that liquidity and other non-default factors can be absorbed into the risk-neutral mean loss rate in pricing defaultable bonds. One simply reinterprets $\widetilde{s}_t$ to include an additive term l_t, possibly state-dependent, that captures the effect of non-default factors. The mean loss rate can still be modeled by an affine function of state variables, preserving the original structure and results.

To re-interpret Equation (15), let us note that the non-default factors can affect expected returns in two ways. First, there is a compensation if they vary with the state variables. This is similar to the compensation for changes in default risk. Second, the modified risk-neutral mean loss rate now incorporates non-default factors, while the physical mean loss rate does not. This difference represents a direct compensation for the non-default factors.

II. ESTIMATION OF DEFAULT EVENT RISK PREMIUM

As should be evident, the issue of expected return is reduced to a standard calculation using estimated model parameters (within the affine class, say) when default event risk is not priced. Since this part of the problem is well understood from the term structure literature, estimation of the defaultable bond expected return boils down to estimation of the default event risk premium.

We focus here on the practical issues associated with this estimation.

Methodology

Consider the yield spread at time t on a zero-coupon bond with maturity T. It is noted that

$$\begin{aligned}
\text{Spread}(t,T) &= -\frac{1}{T-t}\ln\frac{v(t,T)}{p(t,T)} \\
&= -\frac{1}{T-t}\left(\ln\widetilde{E}_t\left(e^{-\int_t^T(r_u+\widetilde{s}_u)du}\right) - \ln\widetilde{E}_t\left(e^{-\int_t^T r_u du}\right)\right) \\
&= -\frac{1}{T-t}\left(\ln\widetilde{E}_t\left(e^{-\int_t^T(r_u+\mu_u s_u+l_u)du}\right) - \ln\widetilde{E}_t\left(e^{-\int_t^T r_u du}\right)\right)
\end{aligned} \tag{16}$$

where $\widetilde{s}$, s, μ, and l are, respectively, the risk-neutral mean loss rate, the physical mean loss rate, the default event risk premium, and the non-default component of the mean loss rate, and $\widetilde{E}_t(\cdot)$ denotes the expectation under the risk-neutral measure.

This expression shows that any empirical feature of the yield spread must be attributable to one of four sources: 1) expected default and recovery through s; 2) risk premium for default event risk through μ; 3) risk premium for variations in the spread through the drift adjustments on $\widetilde{s}$ when changing to the risk-neutral measure; and 4) non-default reasons such as tax and liquidity through l. Abstracting away from non-default sources for the time being, the spread $\widetilde{s}$ and the risk premium for variations in the spread can be identified using standard affine term structure estimation methods. If the physical mean loss rate s can be separately estimated, the default event risk premium μ can be inferred as $\mu = \widetilde{s}/s$.

Risk-Neutral Mean Loss Rate

In the first step of the procedure, the risk-neutral mean loss rate $\widetilde{s}$ would have to be estimated. As a more specific and workable approach, one can model the risk-neutral mean loss rate as an affine function of default-free short rate factors, a "spread index" factor, plus a firm-specific latent factor.

Short rate factors are included to account for the well-documented relationship between yield spread and the default-free term structure (see Duffee [1998]). The spread index can, for example, be the difference between Moody's Baa and Aaa yield indexes, which is meant to capture a major part of the systematic variation in the cross-section of yield spread changes (see Collin-Dufresne, Goldstein, and Martin [2001]). The short rate factors and the spread index factor can be jointly estimated in the first stage, and, given their estimated values, in the second stage the firm-specific factor can be inverted from bond prices.

The short rate factors and the spread index factor should capture enough of the systematic variation in the spreads that there should be no significant remaining risk premium for the latent firm-specific factor. The presence of a firm-specific default risk premium suggests that the inferred firm-specific factor probably subsumes other common factors contributing to changes in the yield spread. This in no way affects subsequent estimation of the default event risk premium. The Jarrow-Lando-Yu framework of diversifiable default risk, for example, can be easily extended to include a firm-specific component in the default intensity, encompassing specifications such as that used by Duffee [1999].

Physical Mean Loss Rate

To simplify estimation of the mean loss rate, one can assume that the recovery, given default, is a constant equal to some measure of a historical average such as those compiled by Moody's. The physical mean loss rate is then equivalent to the physical default intensity. Still, firm-specific estimates of a time-varying physical default intensity are difficult to obtain. Since default is a rare event, such estimates would probably have to be based on a representative firm analysis using financial statement variables and equity returns. This approach has been applied to bankruptcy prediction (see Altman [1968, 1993], Shumway [2001]), but it has not been tested for its ability to produce accurate default probabilities.

An alternative is to use contingent claims analysis to calculate default probabilities based on the time-varying capital structure of a firm (Merton [1974]). This approach is used by KMV to construct an expected default frequency measure, and presumably is also the approach that Moody's uses to rate bond issues. (These are proprietary models employing proprietary data, making them infeasible for academic research.)[8]

If one gives up the hope of using firm-specific and time-varying default intensities, Moody's historical default rates by credit rating class can be used. These are estimated by following a cohort of firms through a given horizon, and then averaging the default rates across all possible starting years in the sample. Or, one can use the average one-year transition matrix to compute default rates with various horizons.

Because of the time aggregation in this procedure, the ability to estimate a random process for μ is lost. As a result, one may simply assume that μ is equal to a constant, and proceed to calculate the implied physical default rates by integrating over the risk-neutral default intensity divided by μ. By minimizing the distance between these implied values and those from Moody's, one obtains an estimate of μ. This can be used to test the diversification condition (whether μ is equal to one) hypothesized by JLY.

Caution must be exercised in interpreting the result obtained by assuming a constant event risk premium. If reality dictates a time-varying one, and especially if conditions are characterized by occasional episodes of high values such as during market crashes, such a simplified approach may lead to erroneous or incomplete conclusions.

For more insight into this issue, let us consider the relation between expected default and the yield spread. Exhibit 1 shows the speculative-grade 12-month trailing

EXHIBIT 1
Speculative-Grade Trailing 12-Month Default Rate and Average Speculative-Grade Spread Against Treasuries

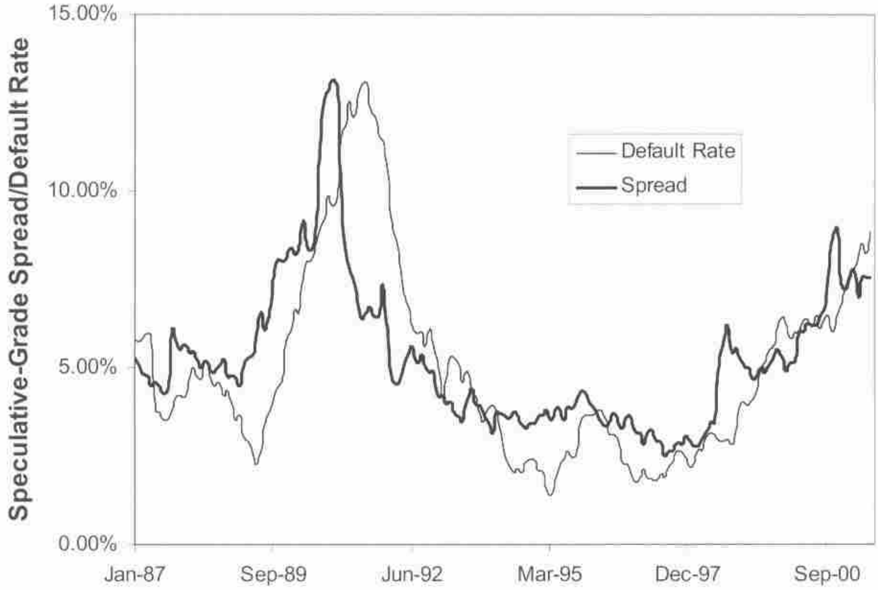

Default rates are obtained from Moody's. Spreads are obtained from Datastream. Sample period is January 1987-August 2001.

default rate as given by Moody's and the speculative-grade spread against Treasuries as given by Datastream. The trailing default rate is calculated by identifying all speculative-grade issuers and then computing the fraction of issuers that have gone into default in the following year.

It is clear from Exhibit 1 that the two series follow the same trend over the sample period. It also appears that the spread leads the default rate by about one year despite the use of a trailing measure. This is because the spread is calculated using bonds whose maturities are most likely beyond one year. If we could use an expected default rate over an appropriate horizon, say, the average duration of bonds in the speculative-grade index, such lead-lag effects will probably be reduced.

Notable in Exhibit 1 are times that dramatic changes in the spread are not accompanied by concurrent or lagged responses in the default rate. The latter half of 1998 seems to be a good example. This episode has been commonly attributed to "low liquidity and shifting investor preferences" (Keenan [2000], p. 7).

If there is no change in the physical mean loss rate, in our framework yield spreads can change for three possible reasons: changes in drift adjustments on the state variables, changes in the event risk premium, and changes in non-default factors such as liquidity. A priori there is no reason to rule out any of these possibilities. Since all three may be intricately linked in practice, the identification of the event risk premium becomes even more daunting.

Non-Default Factors

Estimation of the event risk premium also requires a "clean" risk-neutral mean loss rate that is uncontaminated by taxes, liquidity, and other non-default factors.[9] Although a fully specified model of non-default sources of the spread remains elusive, ad hoc adjustments have been suggested. For example, Jarrow, Lando, and Yu [2001] subtract the average value of the Aaa-implied intensity from the intensities implied from other ratings. They argue that this is a crude way of eliminating the effect of non-default factors.

We first focus on the role of non-default factors in explaining the short-end spread, then revisit the JLY example with another ad hoc but potentially more realistic adjustment for the effect of non-default factors. It is shown that this produces implied default rates that better conform to Moody's estimates.

The purpose of our example is to illustrate the importance of some form of adjustment for non-default factors before one can embark on a more rigorous analysis of the default event risk premium.

EXHIBIT 2
Term Structures of Spreads and Conditional Default Rates

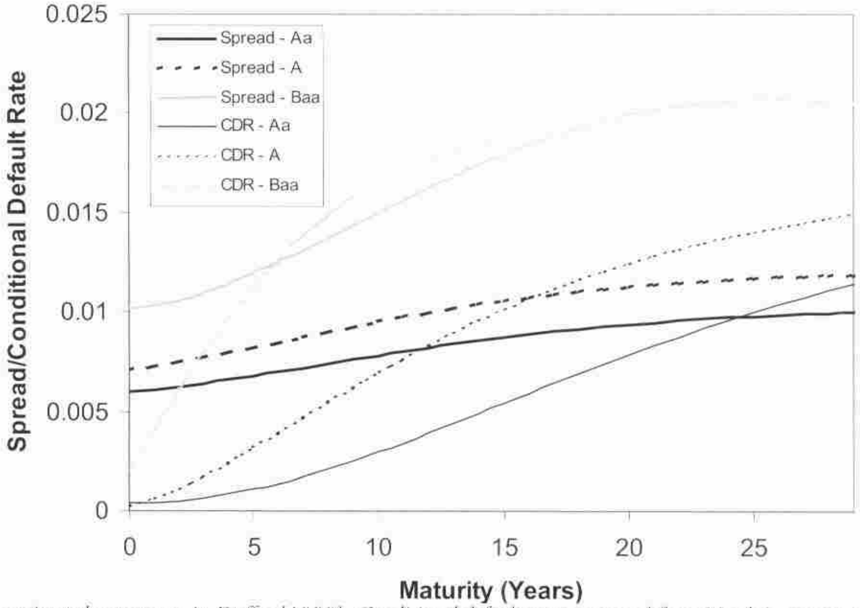

Spreads based on estimated parameters in Duffee [1999]. Conditional default rates computed from Moody's average one-year transition matrix 1980-1999.

Short-End Spreads. Consider the yield spread on a zero-coupon bond as given in Equation (16). Assuming continuity of the processes, it is easy to see that

$$\text{Spread}\,(t, t+) = \widetilde{s}_t \tag{17}$$

On the other hand, the probability of survival until time T is

$$\text{SP}\,(t, T) = E_t\left(e^{-\int_t^T \lambda_u du}\right) \tag{18}$$

where λ_u is the physical default intensity, and $E_t(\cdot)$ denotes the expectation under the physical measure. One can then calculate the instantaneous conditional default rate as

$$\text{CDR}\,(t, T) = -\frac{\partial \ln \text{SP}\,(t, T)}{\partial T} \tag{19}$$

This is the default probability per unit of time, conditional on no default within the interval (t, T). In the limit, as $T \rightarrow t+$, the conditional default rate becomes λ_t. Assuming diversifiable default risk as in JLY, one has $\lambda_t = \widetilde{\lambda}_t$ and:

$$\text{CDR}\,(t, t+) = \frac{\text{Spread}\,(t, t+)}{1 - \delta_t} \tag{20}$$

which results from the definition of $\widetilde{s}_t$. Therefore, under diversifiable default risk and if there were no non-default factors, there is a simple relation that ties down default rates and spreads at the short end.[10]

The data, however, suggest something much more complex. Exhibit 2 presents two plots. The first plot is the term structure of conditional default rates, assuming that the issuer has a given initial credit rating. This is based on Moody's one-year transition matrix averages between 1980 and 1999.[11] The second plot is the term structure of yield spreads computed from the empirical estimates provided in Duffee [1999].

Although the two plots come from time periods that are not entirely comparable, the stark contrast at the short end probably transcends such differences. For the Aa curves, for example, the short-end spread is roughly 60 basis points (bp). With a recovery rate on senior unsecured debt estimated by Moody's to be about 44%, this translates into a default intensity of about 100 bp, compared to the conditional default rate measured at the end of year one of only 4 bp.

This is the most extreme among the three cases presented. Even in the mildest case of Baa-rated issuers, a short-end spread of 100 bp yields an intensity of about 180 bp, much higher than the corresponding short-end default rate of 20 bp. The relation in Equation (20) is therefore grossly violated.

Without non-default-related factors, to reconcile the magnitude of the short-end spread with such a low default rate would require an incredibly high μ_t—much too high, in fact, to be compatible with the apparent agreement in the mid- to long-maturity range.[12] To explain this observation, it seems that one is forced to appeal to non-default-related causes such as liquidity and taxes, whose effect on required yields may decline as a function of maturity. These considerations may also affect expected returns because in Equation (15), the risk-neutral "expected loss" $\widetilde{s}_t$ now includes the effect of liquidity and taxes, while the physical expected loss s_t does not.

Adjustment for Non-Default Factors. To account for the effect of non-default factors, one needs to consider an adjustment to the default intensity. In the absence of any guiding empirical work, it is noted that a constant adjustment to the intensity leads to a constant non-default component of the spread, while adjusting all bond prices downward by a constant fraction leads to a spread adjustment that approaches zero as maturity increases. Specifically, it decreases as $1/T-t$. The realistic situation perhaps lies somewhere between these two cases.

For simplicity, we assume diversifiable default risk so that the example is not complicated by a non-trivial μ_t process. For computation of default rates and spreads at time t, let the adjusted default intensity be defined by

$$\lambda_u' = \lambda_u - \frac{\alpha}{u-t+\beta},\ u \geq t \tag{21}$$

where α and β are positive constants. It is straightforward to show that various adjusted quantities of interest (prime variables) are:

$$\begin{aligned}
\text{Spread}'(t,T) &= \text{Spread}(t,T) - \frac{\alpha(1-\delta)}{T-t}\ln\left(\frac{T-t+\beta}{\beta}\right) \\
\text{SP}'(t,T) &= \text{SP}(t,T)\left(\frac{T-t+\beta}{\beta}\right)^{\alpha} \\
\text{CDR}'(t,T) &= \text{CDR}(t,T) - \frac{\alpha}{T-t+\beta}
\end{aligned} \tag{22}$$

where a constant recovery rate δ is assumed. Hence after adjusting for non-default, "true" spreads are narrower, "true" survival probabilities are greater, and "true" conditional default rates are lower. Focusing on the conditional default rate in Equation (22), one can see that the parameters α and β control both its level at the short end and the speed at which the adjustment term decays with maturity. Clearly, these parameters can be chosen in some suitable fashion to produce a good match of spreads and conditional default rates at the short end in relation to Equation (20).[13]

Term Structure of Implied Conditional Default Rates. The non-default adjustment can shed light on another interesting observation concerning the shape of the term structure of implied conditional default rates. JLY, for example, have relied on notion of asymptotic equivalence between the risk-neutral and the physical default intensity to argue that the term structure of physical default rates can be computed using the risk-neutral intensity extracted from corporate bond prices. They show, however, that the result of such an exercise does not resemble that produced using Moody's one-year rating transition matrix, which is based on historical default data.

Specifically, for investment-grade issuers, the risk-neutral default intensity is high and flat because of the fast mean-reversion of the estimated risk-neutral intensity under the physical measure, while the latter starts out low but gradually increases with maturity. One can experiment with changing the initial value of the risk-neutral default intensity, but this does little to reconcile the differences, because, as Duffee [1999] shows, the intensity is bounded below by a positive constant, and the process mean-reverts rather quickly under the physical measure.

This observation is reproduced in Exhibit 3 with a number of plots for the term structure of conditional default rates. The "unadjusted" series is the result of applying the Jarrow-Lando-Yu methodology to Duffee's intensity estimates for Baa-rated issuers. The "actual" series is the Baa default rate computed from Moody's average one-year transition matrix between 1980 and 1999. Under JLY's diversifiable default risk hypothesis, these two curves ought to be identical. In reality, however, they are indistinguishable only in the mid- to long-maturity range but differ significantly at the short end.

Incorporating the non-default adjustment, the parameters α and β are chosen to minimize the deviation of the non-default-adjusted conditional default rate of Equation (22) from the actual at the short end.[14] Evidently, such a simple model of liquidity, with only two parameters, does a good job in isolating the effect of liquidity and other confounding factors.[15]

EXHIBIT 3
Term Structures of Conditional Default Rates

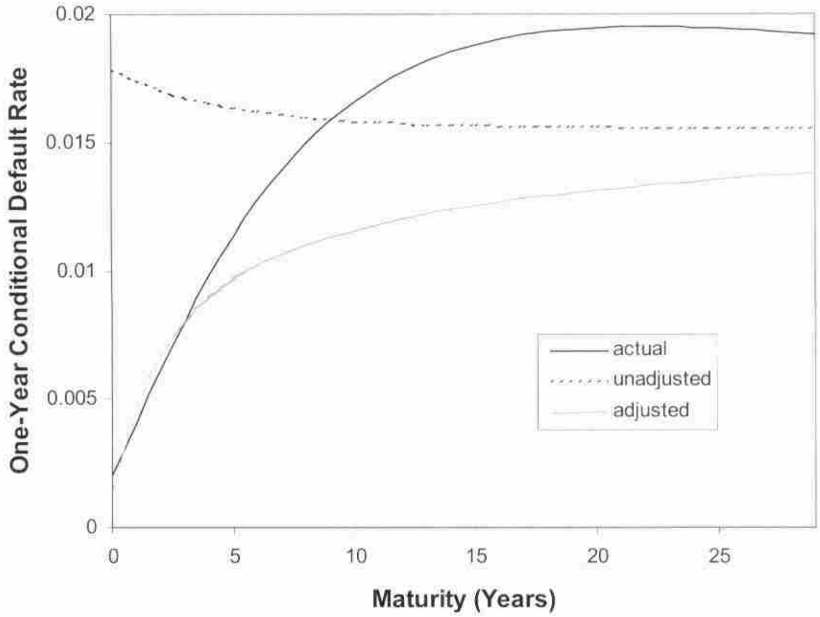

Actual series is Baa default rate computed from Moody's average one-year transition matrix between 1980 and 1999. Unadjusted series is inferred from Duffee's [1999] estimated default intensity for Baa-rated issuers. Unadjusted series is adjusted for liquidity premiums using Equation (21), yielding the adjusted series.

As expected, the difference between the two curves is minimal at the short end. At the long end, their difference could be due to a non-trivial μ_r, but is likely to be statistically insignificant, given the precision of Duffee's estimates. As a result, the overall shape of the term structure of "implied" physical default rates is much closer to the one obtained from actual default data.

What was previously perceived as a slowly mean-reverting intensity process necessary to generate the gently upward-sloping default rates may simply be an artifact of the declining liquidity premium with maturity.

III. UNDERSTANDING EMPIRICAL IRREGULARITIES

Some of the credit risk "irregularities" uncovered in recent empirical studies can be explained by the notion of default event risk and its associated risk premium. Two examples are considered. The first is the finding by Liu, Longstaff, and Mandell [2000] that the excess return on LIBOR bonds as implied by swap-Treasury spreads is negative in over half of their sample period. The second is the finding by Dichev [1998] that firms with higher bankruptcy risk do not earn higher subsequent returns on their stocks.

Credit Premium in the Swaps Market

In a study of the swaps market, Liu, Longstaff, and Mandell [2000] (LLM) compute the expected return on bonds issued by the fictitious "LIBOR" issuer using the intensity implied from swap-Treasury spreads. They produce the puzzling result that the part of this expected return attributable to credit risk—what they call "credit premium"—is negative in over half of their sample period. We provide an explanation of this puzzle.[16]

The interest rate swaps market differs from the U.S. corporate bond market in the source of its credit risk. Corporate bonds are issued by companies, and hence reflect their individual creditworthiness. Interest rate swaps contracts themselves entail minimal credit risk from the counterparties because of contract design.[17] Instead, the credit risk reflected in the swap rates comes from the three-month LIBOR that determines the floating side of the swap, which is agreed upon by a consortium of large international banks located in London. This has led JLY to suspect that the default event underlying the swap-Treasury spread is systematic rather than diversifiable, hence commanding its own risk premium.

Consider the instantaneous expected return on a default-free and a defaultable zero-coupon bond as given in Equations (8) and (15). Taking the difference, the risk premium is given by

$$H_2'\nu'\widetilde{s}_t + \widetilde{s}_t - s_t \tag{23}$$

This shows that the risk premium has two parts. The first part is due to changes in the spread, and the second part is due to the systematic nature of the default event. Since LLM consider only the first part, it is informative to take a detailed look at the entire expression. Using the relation $\widetilde{s}_t = \mu_t s_t + l_t$ where l_t represents the effect of the liquidity differential between swap contracts and Treasury bond benchmarks used in computing the spread, we see that the risk premium is

$$H_2'\nu'\mu_t s_t + (\mu_t - 1)\, s_t + H_2'\nu' l_t + l_t \tag{24}$$

Here, the first term is a compensation for changes in default risk. The second term is a default event risk premium. The third and fourth terms are liquidity premiums; the fourth term represents a direct compensation and the third a compensation for changes in liquidity risk. One would naturally consider the sum of the first two terms to define the credit premium.

LLM, however, define their credit premium as the sum of the first and *the third* terms. They argue that this can turn negative primarily because l_t may be negative (meaning that swap contracts are more liquid than on-the-run Treasury bonds). Consequently, the "liquidity-adjusted" credit premium (essentially just the first term) will be higher, and in fact can become positive in their four-factor Vasicek model.

While the logic of this argument is correct, we see that it ignores an important piece of the puzzle.[18] With a non-trivial μ_t (higher than one on average) affecting the swap market, even LLM's liquidity-adjusted definition will understate the true credit premium. As argued above, this is quite plausible, given the mechanism for determining the underlying LIBOR. Therefore, the credit premium can be positive even in the absence of the alleged swap-Treasury liquidity differential.

Default Risk and Stock Returns

Dichev [1998] examines whether bankruptcy risk is systematic. The approach is to regress realized monthly stock returns, which are proxies for expected return and systematic risk, onto proxies of bankruptcy risk, one of them Altman's [1968, 1993] Z-score. Dichev finds that the regression coefficient of bankruptcy risk is not significantly different from zero, and is even negative in some sample periods, suggesting that firms with higher bankruptcy risk earn lower-than-average expected returns. He concludes that default risk is not systematic.

The expected return decomposition of Equation (24) suggests additional complexity in answering this question.[19] After all, the default likelihood under the physical measure, which Altman's Z-score purports to reflect, is only one of several determinants of corporate bond (as well as common stock) expected returns. It is therefore not difficult to give a number of conjectures that might explain Dichev's empirical finding.

First, note that Dichev's data include delisting returns. To the extent that the monthly returns of every firm in the sample, up to and including its potential bankruptcy, are correctly accounted for by this procedure, Equation (24) would be the appropriate expression for assessing expected return.

Shumway [1997] and Shumway and Warther [1999], however, show that this is not the case. Not including or inadequately accounting for the high negative returns on defaulted firms in the analysis would amount to having an expected return equal to just the first two terms in Equation (23). Intuitively, when the return on bankrupt firms is overlooked, realized return will tend to be overstated. This "survival bias" is more severe for speculative-grade issuers since their default probabilities are higher than those of investment-grade issuers. Unfortunately, it is easy to see that this line of thinking cannot explain why firms with high bankruptcy risk tend to have below-average returns.

Doubts about the data aside, a plausible explanation for Dichev's findings is that speculative-grade default and recovery rates may be less sensitive to the state variable than these elements for investment-grade counterparts. This will generate a difference in the market price of risk parameter ν' that might offset a positive relation between bankruptcy risk and expected return.

For example, consider the distinction between an oil company heavily invested in oil exploration and, say, General Motors. The oil company's fortune depends a great deal on the success of its small number of current projects, which are highly risky but idiosyncratic. GM is a much safer bet, perhaps, but automobile sales are typically highly procyclical. Note, however, that this argument applies in a one-factor setting. When multiple factors drive changes in default intensity over the business cycle, various factor exposures can also produce offsetting effects on expected return when moving from high to low credit quality, thus explaining the lack of correlation between bankruptcy probability and expected return.

A second explanation is that speculative-grade issuers may command a smaller event risk premium than investment-grade issuers. This has a similar offsetting effect on expected returns. This explanation is plausible if, for example, low-quality or junk issues tend to be small and to be held in large portfolios, while high-quality issues tend to be large and held by only a small number of investors (which implies that their portfolios are rather concentrated). Alternatively, investors in junk bonds may be used to the notion of default risk, and thus may be risk-neutral with respect to default events. In comparison, defaults by investment-grade bonds are rare, and to that extent investors in high-quality debt may well demand compensation for a surprising default event.

Without actual empirical work, these explanations are merely speculations—although they show that Dichev's conclusion is limited by his definition of "systematic default risk." Before one accepts any non-risk-based explanation of Dichev's results, it would seem necessary to achieve a more detailed understanding of bankruptcy risk than is currently available. Our work on the structure of expected return helps to do this.

IV. CONCLUSION

With the proliferation of standard term structure modeling techniques in the credit risk literature, one must exercise caution in translating certain procedures for default-free bonds to defaultable bonds. We emphasize expected return calculations as an example.

Using a simple affine framework, we show that the expected return on defaultable bonds can be decomposed into a default-free component, a compensation for variations in default risk, and a compensation for bearing the default event. The last component has heretofore been ignored in the empirical credit risk literature. It originates from a "survival bias" that arises, for example, whenever one drops defaulted firms in computing realized returns on a bond portfolio.

The expected return decomposition is then used in an attempt to resolve several empirical puzzles in the credit risk literature. First, the default event risk premium can help to explain the negative credit premium identified by Liu, Longstaff, and Mandell [2000] in their study of the LIBOR swap-Treasury spread. Second, the default event risk premium may be one of the reasons that Dichev [1998] finds an insignificant relation between stock return and default risk.

We also present a general methodology for estimating the default event risk premium, although many difficulties remain for this estimation, including the specification and identification of stochastic processes describing the physical default intensity as well as the non-default part of the yield spread.

ENDNOTES

The author thanks for helpful comments Magnus Dahlquist, Darrell Duffie, Philippe Jorion, and Haitao Li; and participants at the 12th Annual Derivative Securities Conference, the International Credit Risk Conference at HEC Montreal, and the Credit Risk Conference at the Stockholm Institute for Financial Research.

[1]For affine term structure models, see Duffie and Kan [1996]. Dai and Singleton [2000] present a detailed classification and empirical study of affine models. Duarte [2000] and Duffee [2002] extend the standard affine framework using more general specifications of the market price of risk. The latter derives general expressions of instantaneous expected return on zero-coupon bonds in affine models and their extensions.

[2]Intuitively, this is not unlike the joint modeling of stock and option prices in order to identify the risk premium on jump factors commonly used to explain the volatility smile.

[3]For a more general specification that allows such dependence, see Duffee [1999].

[4]For simplicity, other contributors of the spread, such as liquidity, taxes, and asymmetric information, are ignored. Considerations for these factors can be added later on.

[5]The existence of a point process intensity under equivalent changes of measure is shown by Artzner and Delbaen [1995].

[6]A debate on the importance of the survival bias for equity returns can be found in Brown, Goetzmann, and Ross [1995] and Li and Xu [2000].

[7]This term includes the usual interest rate risk premiums.

[8]Vassalou and Xing [2002] use Merton's [1974] model to compute default probabilities. They show that this measure of default risk works better than accounting-based measures or default spreads in explaining the cross-section of stock returns. Thus it appears that although the structural approach runs into difficulties as a valuation model for corporate debt it may be well-suited for constructing physical default probabilities (see Jones, Mason, and Rosenfeld [1984]).

[9]See Huang and Huang [2000], Duffie and Lando [2001], Elton et al. [2001], and Schultz [2001] for discussion of these issues.

[10]In deriving this result, it is assumed that the intensities are driven by diffusion state variables, but the result generalizes to jump-state variables as long as the rate of jumps is finite.

[11]With the one-year transition matrix, only yearly estimates of survival probabilities can be computed. Therefore, standing in year 0, it is assumed that the probability of surviving until year n is p_n. The conditional default rate in year n is then computed as $(p_n - p_{n+1})/p_n$. This is of course a discrete-time approximation of the instantaneous conditional default rate as defined before.

[12]The same is true for ν', the drift adjustment of the risk-neutral expected loss rate. Jarrow, Lando, and Yu [2001] provide a numerical example illustrating this point.

[13]JLY have argued that by adjusting for the effect of state taxes and liquidity as in Elton et al. [2001] and others, short-end spreads can be boiled down to physical mean loss rates alone. Ericsson and Renault [1999] construct a structural model of liquidity and credit risk. One of their findings is that the liquidity component of the spread is a decreasing function of bond maturity, and that the price discount approaches a constant as maturity increases. These results support the interpretation of the ad hoc adjustment given above as due to non-default factors. While the simple adjustment here ignores the interaction between credit and liquidity risk that is the focus of Ericsson and Renault [1999], this can be incorporated into the general intensity-based framework by assuming state-dependent liquidity adjustments.

[14]Specifically, it minimizes the sum of squared difference between the actual and adjusted one-year conditional default rate over $T - t = 0, 1, \ldots, 4$. The long end is left out of the optimization since liquidity likely has less of a role there.

[15]Ericsson and Renault [1999] show that the downward-sloping term structure of liquidity spreads can explain the apparent failure of structural models to produce a flat term structure of yield spreads for investment-grade issuers. Theirs is another look at essentially the same problem from a different angle.

[16]LLM use a four-factor Vasicek model to capture features of the swap-Treasury spread. In our illustration, we use a significantly simpler setup, which incidentally does not allow the possibility of a negative credit premium (although the basic intuition easily carries over to the more complex model).

[17]For example, see Duffie and Huang [1996], Jarrow and Turnbull [1997], Li [1998], and Collin-Dufresne and Solnik [2001].

[18]On a minor note, the argument for a negative liquidity premium suggests that the default part of the spread should be measured with respect to a fictitious Treasury bond with the same liquidity as a swap contract. This seems to conflict with LLM's earlier choice of the specialness premium in the Treasury market as the liquidity component of the swap spread process.

[19]Equation (23) addresses the expected return on defaultable bonds, while Dichev [1998] studies the relation between bankruptcy risk and stock return. Yet it is not unreasonable to assume that the insights from studying the first problem will have some bearing on the second.

REFERENCES

Altman, E.I. *Corporate Financial Distress and Bankruptcy: A Complete Guide to Predicting and Avoiding Distress and Profiting from Bankruptcy*. New York: John Wiley & Sons, 1993.

——. "Financial Ratios, Discriminant Analysis, and the Prediction of Corporate Bankruptcy." *Journal of Finance*, 23 (1968), pp. 589-609.

Artzner, P., and F. Delbaen. "Default Risk Insurance and Incomplete Markets." *Mathematical Finance*, 5 (1995), pp. 187-195.

Bakshi, G., D. Madan, and F. Zhang. "Investigating the Sources of Default Risk: Lessons from Empirically Evaluating Credit Risk Models." Working paper, University of Maryland, 2001.

Brown, S., W. Goetzmann, and S.A. Ross. "Survival." *Journal of Finance*, 50 (1995), pp. 853-873.

Collin-Dufresne, P., R. Goldstein, and S. Martin. "The Determinants of Yield Spread Changes." *Journal of Finance*, 56 (2001), pp. 2177-2207.

Collin-Dufresne, P., and B. Solnik. "On the Term Structure of Default Premia in the Swap and LIBOR Markets." *Journal of Finance*, 56 (2001), pp. 1095-2015.

Cox, J., J. Ingersoll, and S. Ross. "A Theory of the Term Structure of Interest Rates." *Econometrica*, 53 (1985), pp. 385-407.

Dai, Q., and K. Singleton. "Specification Analysis of Affine Term Structure Models." *Journal of Finance*, 55 (2000), pp. 1943-1978.

Dichev, I. "Is the Risk of Bankruptcy a Systematic Risk?" *Journal of Finance*, 53 (1998), pp. 1131-1148.

Duarte, J. "The Relevance of the Price of Risk in Affine Term Structure Models." Working paper, University of Chicago, 2000.

Duffee, G.R. "Estimating the Price of Default Risk." *Review of Financial Studies*, 12 (1999), pp. 197-226.

——. "The Relation between Treasury Yields and Corporate Bond Yield Spreads." *Journal of Finance*, 53 (1998), pp. 2225-2242.

——. "Term Premia and Interest Rate Forecasts in Affine Models." *Journal of Finance*, 57 (2002), pp. 405-443.

Duffie, J.D., and M. Huang. "Swap Rate and Credit Quality." *Journal of Finance*, 51 (1996), pp. 921-949.

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

Duffie, D., and D. Lando. "Term Structures of Credit Spreads with Incomplete Accounting Information." *Econometrica*, 69 (2001), pp. 633-664.

Duffie, J.D., L. Pedersen, and K.J. Singleton. "Modeling Sovereign Yield Spreads: A Case Study of Russian Debt." Working paper, Stanford University, 2001.

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

——. "Modeling Term Structures of Defaultable Bonds." *Review of Financial Studies*, 12 (1999), pp. 687-720.

Elton, E., M. Gruber, D. Agrawal, and C. Mann. "Explaining the Rate Spread on Corporate Bonds." *Journal of Finance*, 56 (2001), pp. 247-278.

Ericsson, J., and O. Renault. "Credit and Liquidity Risk." Working paper, McGill University, 1999.

Huang, J., and M. Huang. "How Much of the Corporate-Treasury Yield Spread is Due to Credit Risk?" Working paper, The Pennsylvania State University, 2000.

Jarrow, R.A., D. Lando, and F. Yu. "Default Risk and Diversification: Theory and Applications." Working paper, UC-Irvine, 2001.

Jarrow, R.A., and S.M.Turnbull. "Pricing Derivatives on Financial Securities Subject to Credit Risk." *Journal of Finance*, 50 (1995), pp. 53-85.

——. "When Swaps Are Dropped." *Risk*, Vol. 10, No. 5 (1997), pp. 70-75.

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

Keenan, S.C. "Historical Default Rates of Corporate Bond Issuers: 1920-1999." Moody's, 2000.

Lando, D. "On Cox Processes and Credit Risky Securities." *Review of Derivatives Research*, 2 (1998), pp. 99-120.

Li, H. "Pricing of Swaps with Default Risk." *Review of Derivatives Research*, 2 (1998), pp. 231-250.

Li, H., and Y. Xu. "Can Survival Bias Explain the Equity Premium Puzzle?" Working paper, Cornell University, 2000.

Liu, J., F. Longstaff, and R. Mandell. "The Market Price of Credit Risk: An Empirical Analysis of Interest Rate Swap Spreads." Working paper, UCLA, 2000.

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

Schultz, P. "Corporate Bond Trading Costs: A Peek Behind the Curtain." *Journal of Finance*, 56 (2001), pp. 677-698.

Shumway, T. "The Delisting Bias in CRSP Data." *Journal of Finance*, 52 (1997), pp. 327-340.

——. "Forecasting Bankruptcy More Accurately: A Simple Hazard Model." *Journal of Business*, 74 (2001), pp. 101-124.

Shumway, T., and V. Warther. "The Delisting Bias in CRSP's Nasdaq Data and its Implications for the Size Effect." *Journal of Finance*, 54 (1999), pp. 2361-2379.

Vassalou, M., and Y. Xing. "Default Risk in Equity Returns," Working paper, Columbia University, 2002.

To order reprints of this article please contact Ajani Malik at amalik@iijournals.com or 212-224-3205.